

Applications of hidden Markov models in neutron star astronomy

Liam Michael Dunn

ORCID: [0000-0002-1769-6097](https://orcid.org/0000-0002-1769-6097)

August, 2024

A thesis submitted in complete fulfillment
of the requirements for the degree of

Doctor of Philosophy
Physics
The University of Melbourne

© Liam Michael Dunn
All rights reserved, 2024

To my parents.

...In that Empire, the Art of Cartography attained such Perfection that the map of a single Province occupied the entirety of a City, and the map of the Empire, the entirety of a Province. In time, those Unconscionable Maps no longer satisfied, and the Cartographers Guilds struck a Map of the Empire whose size was that of the Empire, and which coincided point for point with it. The following Generations, who were not so fond of the Study of Cartography as their Forebears had been, saw that that vast Map was Useless, and not without some Pitilessness was it, that they delivered it up to the Inclemencies of Sun and Winters. In the Deserts of the West, still today, there are Tattered Ruins of that Map, inhabited by Animals and Beggars; in all the Land there is no other Relic of the Disciplines of Geography.

— Suárez Miranda, *Viajes de varones prudentes*, Libro IV, Cap. XLV, Lérida, 1658

Abstract

Since their discovery nearly 60 years ago, neutron stars have proved to be a rich source of astrophysical insight. Most often they are observed as *pulsars* — rapidly rotating neutron stars which emit beams of electromagnetic radiation aligned with their magnetic axis, sweeping past the Earth and producing a sequence of pulses at observatories operating across the electromagnetic spectrum. Their rotation is remarkably stable, and this allows for detailed study of the subtle variations in the observed pulse train, an undertaking broadly known as *pulsar timing*. Neutron stars are also a promising source of persistent gravitational radiation (*continuous gravitational waves*, or CW) in the frequency band accessible to the current generation of terrestrial interferometric gravitational wave detectors such as the Laser Interferometric Gravitational-Wave Observatory (LIGO), Virgo, and the Kamioka Gravitational Wave Detector (KAGRA). At the time of writing, no detection of CW emission has been made.

Both pulsar timing and searches for CW emission from neutron stars are demanding from a data analysis perspective, and a wide variety of techniques have been developed to tackle various questions in these areas. This thesis is primarily concerned with the application of hidden Markov models (HMMs) — a particular class of models for stochastic dynamical systems — to both pulsar timing and CW searches.

The application of HMMs to pulsar timing discussed in this thesis is an automated detector of glitches — small instantaneous jumps in a pulsar’s spin frequency. Automatic detection of glitches is challenging, as any detection scheme must also contend with spin wandering, referred to as timing noise. The HMM glitch detector provides a framework for modelling the timing noise along with any glitches in the data, and a quantitative means of assessing the sensitivity of the detector, which can be translated to upper limits on the size of undetected glitches in a given data set. We present four analyses using this method. The first is an analysis of 282 pulsars released as part of the first UTMOST data release. We find no previously undetected glitches, and set 90% upper limits on the fractional glitch size, denoted $\Delta\nu^{90\%}/\nu$, for each pulsars, with a mean upper limit of 1.9×10^{-8} . The second search is an analysis of 24 years of high-cadence observations of the Vela pulsar from the Mount Pleasant Observatory. The Vela pulsar is noted for its quasiperiodic large glitches, but is also known to exhibit smaller glitches. Here we use the HMM to search for small glitches which were previously missed. We detect one new glitch, with fractional size $\Delta\nu/\nu = (8.19 \pm 0.04) \times 10^{-10}$, and set upper limits on the size of other missed glitches, with a median upper limit of $\Delta\nu^{90\%}/\nu = 1.35 \times 10^{-9}$. The other two searches incorporate data from the recent UTMOST-NS pulsar timing campaign, which operated from 2021 to 2023. One is an “online” search which incorporated new UTMOST-NS observations soon after they were recorded and generated alerts when glitch candidates were detected. Three large glitches were detected by this pipeline,

in PSRs J0835–4510, J0742–2822, and J1740–3015. Alerts were produced for all three glitches in the days following the event by both the UTMOST collaboration and a number of other observatories. The other is an “offline” search which was performed after the UTMOST-NS programme finished and incorporates both UTMOST-NS observations and observations from the original UTMOST programme, covering 163 pulsars in total. No additional glitches are detected in the offline search, but thanks to explicit incorporation of measurements of the timing noise in each pulsar we are able to set more constraining upper limits than the previous analysis of UTMOST data alone, with a mean upper limit of $\Delta\nu^{90\%}/\nu = 5.3 \times 10^{-9}$.

We also discuss aspects of pulsar timing beyond glitch detection. We detail the effect of periodic observation scheduling on the measurement of glitch sizes, and show that under realistic observing scenarios, glitch sizes can be misestimated by an order of magnitude or more. We present a case study of a glitch in PSR J1709–4429 which was originally estimated to have a size of $\Delta\nu/\nu = (54.6 \pm 1.0) \times 10^{-9}$, and give a revised estimate of $\Delta\nu/\nu = (2432.2 \pm 0.1) \times 10^{-9}$. We show that standard pulsar timing techniques are liable to miss this degeneracy, whereas the HMM pulsar timing framework alerts the analyst to the effect via a multimodal posterior on the glitch amplitude.

Additionally, we investigate the characteristics of timing noise in the sample of pulsars observed by UTMOST-NS. Assuming that timing noise strength σ_{RN} scales as $\nu^a |\dot{\nu}|^b$, we find that $a = -0.92^{+0.31}_{-0.29}$, $b = 0.80^{+0.12}_{-0.13}$, i.e. σ_{RN} is approximately correlated with the inverse characteristic age $\tau_c^{-1} \sim \nu^{-1} |\dot{\nu}|$. We also find that measurements of timing noise power spectral density (PSD) appear to be biased towards shallower spectral indices when the observing timespan is increased. Finally, we discuss the interplay between timing noise and measurement of second frequency derivatives, and show that the measurement of $\ddot{\nu}$ is strongly dependent on both the statistical nature of the timing noise and the details of the timing noise model. For example, the uncertainty on $\ddot{\nu}$ is underestimated by a factor of roughly three in cases where the timing residuals PSD has a spectral index of approximately six and the timing noise model does not include modes with period longer than the observing timespan. We verify a predictive theoretical scaling between the magnitude of the measured $\ddot{\nu}$ and a combination of the timing noise strength and observing timespan.

CW searches often contend with large parameter spaces. This problem is exacerbated when spin wandering is taken into account – just as the electromagnetic signal from a pulsar displays long-term stochastic behaviour (timing noise), so too might a CW signal. The value of HMMs in CW analyses lies in their ability to efficiently cover wide parameter spaces and account for spin wandering. In this thesis we present an analysis in which an HMM is used to search for CW emission from neutron stars in five nearby globular clusters, using data from LIGO’s third observing run. The HMM is set up to track not only the frequency of the CW signal over time, but also the phase. Frequency-tracking HMMs have been used in many CW searches previously, but this is the first time that the phase-tracking extension has been applied to a search using real data. No significant candidates are detected, and the strain sensitivity at 95% confidence, $h_{0,\text{eff}}^{95\%}$ is estimated, along with corresponding limits on the neutron star ellipticity $\epsilon^{95\%}$. The best limits are achieved for the cluster NGC 6544 at 211 Hz, with $h_{0,\text{eff}}^{95\%} = 2.7 \times 10^{-26}$, and for the cluster NGC 6397 at 800 Hz, with $\epsilon^{95\%} = 1.6 \times 10^{-7}$.

Declaration

This is to certify that:

- This thesis contains only original work towards a Doctor of Philosophy, except where indicated in the preface.
- Due acknowledgement has been made in the text to all other material used.
- This thesis is fewer than 100 000 words in length, exclusive of tables, figures, bibliographies, and appendices.

Liam Michael Dunn

Preface

This thesis is the original work of the author, describing work done alone and in collaboration with others. This section outlines the contributions by the author and collaborators to each chapter.

Chapter 1 is a literature review written by the author. Thesis supervisor Melatos provided editorial input.

Chapter 2 is published as

L. Dunn, M. E. Lower and A. Melatos (2021), “Effects of periodicity in observations scheduling on parameter estimation of pulsar glitches”, *Monthly Notices of the Royal Astronomical Society* 504.3, pp. 3399–3411.

This work was written primarily by the author. Co-authors provided scientific input and commented on and edited the text. All figures and tables are the work of the author.

Chapter 3 is published as

L. Dunn et al. (2022), “Systematic upper limits on the size of missing pulsar glitches in the first UTMOST open data release”, *Monthly Notices of the Royal Astronomical Society* 512.1, pp. 1469–1482.

This work was written primarily by the author. Co-authors provided assistance with the methodology and with data handling, and provided scientific input and commented on and edited the text. All figures and tables are the work of the author.

Chapter 4 is published as

L. Dunn, A. Melatos, C. M. Espinoza, D. Antonopoulou, and R. Dodson (2023), “A new small glitch in the Vela pulsar discovered with a hidden Markov model”, *Monthly Notices of the Royal Astronomical Society* 522.4, pp. 5469–5478.

This work was written primarily by the author. Co-authors provided assistance with data provision and analysis, provided scientific input, and commented on and edited the text. All figures and tables are the work of the author.

Chapter 5 has been submitted to *Monthly Notices of the Royal Astronomical Society* as

L. Dunn et al. (2024), “First results from the UTMOST-NS pulsar timing programme”.

This work was written primarily by the author. Co-authors were responsible for operating the programme which produced the data analysed in this work, and drafted text in Section 5.2 describing this programme. They also provided general scientific input and commented on and edited the text. Figure 5.1 was produced by Chris Flynn. All other figures and tables are the work of the author.

Chapter 6 has been submitted to Physical Review D for publication as

L. Dunn et al. (2024), “Search for continuous gravitational waves from neutron stars in five globular clusters with a phase-tracking hidden Markov model in the third LIGO observing run”.

This work was written primarily by the author. Co-authors provided assistance with the methodology and provided scientific input and commented on and edited the text. All figures and tables are the work of the author.

Chapter 7 summarizes the thesis and discusses possible avenues of future work, written by the author. Thesis supervisor Melatos provided editorial input.

During the course of their PhD candidature, the author also contributed to other publications which are not included for examination in this thesis but are listed here for completeness.

- Beniwal et al. (2021), “Search for continuous gravitational waves from ten H.E.S.S. sources using a hidden Markov model”, Physical Review D 103.8, 083009.
The author provided advice on methodology.
- Wette et al. (2021), “Deep exploration of continuous gravitational waves at 171–172 Hz in LIGO second observing run data”, Physical Review D 103.8, 083020.
The author developed graphics processing unit-accelerated analysis software and commented on the text.
- Lower et al. (2021), “The impact of glitches on young pulsar rotational evolution”, Monthly Notices of the Royal Astronomical Society 508.3, pp. 3251–3274.
The author performed glitch upper limit analyses, provided scientific input and commented on the text.
- Dunn et al. (2022c), “Graphics processing unit implementation of the \mathcal{F} -statistic for continuous gravitational wave searches”, Classical and Quantum Gravity 39.4, 045003.
The author developed the software described in this publication, performed analyses, drafted text and produced all figures and tables.
- Millhouse et al. (2022), “An updated glitch rate law inferred from radio pulsars”, Monthly Notices of the Royal Astronomical Society 511.3, pp. 3304–3319.
The author provided glitch upper limit data and commented on text.
- Beniwal et al. (2022), “Search for continuous gravitational waves from HESS J1427–608 with a hidden Markov model”, Physical Review D 106.10, 103018.
The author provided advice on methodology.

- Jones et al. (2022), “Validating continuous gravitational-wave candidates from a semicoherent search using Doppler modulation and an effective point spread function”, *Physical Review D* 106.12, 123011.
The author provided scientific input and advice on methodology.
- Lower et al. (2023a), “Rotational and radio emission properties of PSR J0738–4042 over half a century”, *Monthly Notices of the Royal Astronomical Society* 524.4, pp. 5904–5917.
The author shared UTMOST-NS data which was used in the analysis.
- Lower et al. (2023b), “The 2022 High-energy Outburst and Radio Disappearing Act of the Magnetar 1E 1547.0–5408”, *The Astrophysical Journal* 945.2, 153.
The author performed pulsar timing analysis, provided scientific input and commented on text.
- Knee et al. (2024), “Search for continuous gravitational waves directed at sub-threshold radiometer candidates in O3 LIGO data”, *Physical Review D* 109.6, 062008.
The author provided scientific input and advice on methodology, and commented on text.
- Mandlik et al. (2024), “The UTMOST-NS: A fully digital, wide-field transient facility operating at a centre frequency of 831 MHz”, *Monthly Notices of the Royal Astronomical Society* 532.2, pp. 2644–2656.
The author participated in the UTMOST-NS pulsar timing programme, and provided figure 7.

The author is supported by an Australian Government Research Training Program Scholarship, the Rowden White Scholarship, the Tyndall scholarship, the Dieul-Kurzweil scholarship, and funds from the Australian Research Council (ARC) Centre of Excellence for Gravitational Wave Discovery (OzGrav) (project numbers CE170100004 and CE230100016).

Acknowledgments

This thesis owes its existence to the hard work and kindness of many people. First, my thanks to Andrew, not only for his supervision throughout the many ups and downs of the past four years but for taking me on as an undergraduate, which set everything in motion. It's been a lot of fun. Thanks also to the other members of my supervisory panel, Katie and Harry, for their advice throughout.

I am deeply indebted to all of my co-authors, who made the work presented here possible, and especially to Sofia Suvorova, whose code formed the basis for much of this thesis. Special thanks also to the UTMOST team, for their excellent work in operating Molonglo up to the bitter end, and for giving me the chance to participate in the UTMOST-NS project, which has been one of the highlights of this PhD.

Thanks to the various incarnations of the Viterbi and NS/GW groups which I have had the pleasure to be a part of over the years, as well as the wider Astro group.

Thanks to my friends, for whom I am always grateful. And, finally, thank you to Mum, Dad, and Kate, for all your love and support over the years, and everything else besides.

Contents

| | |
|---|-------------|
| Abstract | v |
| Declaration | vii |
| Preface | viii |
| Acknowledgments | xi |
| List of figures | xv |
| List of tables | xxiv |
| 1 Introduction | 1 |
| 1.1 Neutron stars | 1 |
| 1.2 Radio pulsar timing | 6 |
| 1.3 Pulsar spin down | 10 |
| 1.4 Timing noise | 12 |
| 1.5 Glitches | 18 |
| 1.6 Gravitational waves in general | 24 |
| 1.7 Continuous gravitational waves | 27 |
| 1.8 Hidden Markov models | 33 |
| 1.9 Hidden Markov models and continuous waves | 34 |
| 1.10 Hidden Markov models and pulsar timing | 37 |
| 1.11 Thesis outline | 40 |
| 2 Effects of periodicity in observation scheduling on parameter estimation of pulsar glitches | 41 |
| 2.1 Introduction | 42 |
| 2.2 Phase ambiguity | 43 |
| 2.3 Glitches | 48 |
| 2.4 Glitch parameter estimation | 50 |
| 2.5 Re-assessing UTMOST glitches | 59 |
| 2.6 Conclusion | 65 |
| Appendix 2.A1 HMM recipe and parameters | 67 |
| 3 Systematic upper limits on the size of missing pulsar glitches in the first UTMOST open data release | 71 |
| 3.1 Introduction | 72 |

| | | |
|----------|---|------------|
| 3.2 | Data | 73 |
| 3.3 | Hidden Markov model | 73 |
| 3.4 | UTMOST glitches | 79 |
| 3.5 | Size upper limits | 84 |
| 3.6 | Conclusion | 89 |
| | Appendix 3.A1 Follow-up analysis of vetoed candidates | 91 |
| | Appendix 3.A2 Synthetic dataset generation and upper limit estimation | 92 |
| | Appendix 3.A3 Posterior distributions and frequency tracks for HMM analyses | 94 |
| 4 | A new small glitch in Vela discovered with a hidden Markov model | 98 |
| 4.1 | Introduction | 99 |
| 4.2 | HMM glitch detector | 100 |
| 4.3 | A small glitch at MJD 48636 | 103 |
| 4.4 | Other small glitches | 108 |
| 4.5 | Undetected glitches | 111 |
| 4.6 | Conclusion | 114 |
| | Appendix 4.A1 MJD 48636 event as a candidate in previous searches | 115 |
| 5 | First results from the UTMOST-NS pulsar timing programme | 116 |
| 5.1 | Introduction | 117 |
| 5.2 | Observations | 118 |
| 5.3 | Pulsar timing | 123 |
| 5.4 | Timing noise | 127 |
| 5.5 | Second frequency derivatives | 133 |
| 5.6 | Glitches | 141 |
| 5.7 | Conclusion | 155 |
| | Appendix 5.A1 Flux density calibration and telescope sensitivity for the survey period | 158 |
| | Appendix 5.A2 Influence of timing noise on the measured $\dot{\nu}$ for $ \beta - 4 < 1$ | 160 |
| | Appendix 5.A3 HMM setup for offline glitch searches | 161 |
| 6 | Search for continuous gravitational waves from neutron stars in five globular clusters with a phase-tracking hidden Markov model in the third LIGO observing run | 163 |
| 6.1 | Introduction | 164 |
| 6.2 | Algorithms | 165 |
| 6.3 | Setting up a globular cluster search | 171 |
| 6.4 | Candidates | 183 |
| 6.5 | Sensitivity | 187 |
| 6.6 | Conclusion | 190 |
| | Appendix 6.A1 Transition probabilities | 191 |
| | Appendix 6.A2 Detection thresholds | 192 |
| | Appendix 6.A3 Single-interferometer veto | 193 |
| | Appendix 6.A4 Unknown lines | 193 |
| 7 | Conclusions and outlook | 198 |
| 7.1 | Summary | 198 |

| | | |
|---|--|------------|
| 7.2 | Future work: radio pulsars | 202 |
| 7.3 | Future work: continuous gravitational waves | 206 |
| Appendix A Bayesian pulsar timing | | 210 |
| Appendix B A signal consistency test for continuous wave signals | | 212 |
| B.1 | Vetoed and signal consistency tests | 212 |
| B.2 | Doubling coherence time as a signal consistency test | 213 |
| B.3 | Application to candidates from superseded GC search | 214 |
| References | | 223 |

List of figures

| | | |
|-----|---|----|
| 1.1 | P - \dot{P} diagram of the full pulsar population generated from the ATNF pulsar catalogue (Manchester et al., 2005) using PSRQPY (Pitkin, 2018). Pulsars in binaries are circled in red. Lines of constant characteristic age ($\tau_c = P/2\dot{P}$) and surface magnetic field [$B_{\text{surf}} = 3.2 \times 10^{19} \text{ G } (P/1 \text{ s})^{1/2} \dot{P}^{1/2}$] are shown as dashed and dotted lines respectively. The shaded region in the bottom right is below the “death line” as defined in Equation 4 of Zhang et al. (2000), approximately where pulsar radio emission is expected to cease. | 5 |
| 1.2 | Timing residuals for the pulsar PSR J1224–6407, using observations from the UTMOST pulsar timing programme (see Jankowski et al. 2019 and Chapter 5 of this thesis). Significant long-term wandering behaviour (timing noise) is evident. | 12 |
| 1.3 | Scatter plots showing the relationship between the rms timing residuals δt_{rms} and inferred quantities related to the timing noise model (<i>top row</i>) and the spin evolution of the pulsar (<i>bottom row</i>), for 597 pulsars timed as part of the MeerKAT Thousand Pulsar Array project (Keith et al., 2024), as discussed in Section 1.4.2. The dashed line in the top left panel indicates $\delta t_{\text{rms}} = \sigma_{\text{RN}}$. The data are divided according to the value of A_{red} , with $A_{\text{red}} > 10^{-11}$ indicated by the blue points and $A_{\text{red}} < 10^{-11}$ by the orange. Spearman correlation coefficients between δt_{rms} and the other parameters are given in Table 1.1. | 16 |
| 1.4 | Histogram of glitch sizes recorded in the JBO glitch catalogue (Basu et al., 2022). | 19 |
| 1.5 | Timing residuals for the pulsar PSR J0908–4913, using observations from the UTMOST pulsar timing programme. A glitch is present in the data with size $\Delta f \approx 2 \times 10^{-7} \text{ Hz}$, but is not included in the timing model. A clear cuspy feature is visible at the location of the glitch. | 20 |
| 1.6 | Distribution of minimum observed fractional glitch amplitude among pulsars with at least two glitches recorded. The dotted lines indicate the minimum (2.4×10^{-12}) and maximum (1.3×10^{-7}) 90% confidence upper limits on the fractional size of undetected glitches derived from the analyses described in Chapter 5, while the solid line shows the median 90% upper limit (2.1×10^{-9}). | 21 |
| 1.7 | Representative noise amplitude spectral density curves for the Hanford (H1) LIGO detector during the first observing run (O1) (Sigg, 2016) and second half of the third observing run (O3b) (Goetz, 2021). | 26 |

| | | |
|-----|--|----|
| 2.1 | Mean fractional scheduling aperiodicity $\langle \epsilon_i \rangle$ as a function of scheduling quasiperiod T for the UTMOST observations of PSR J1452–6036. | 45 |
| 2.2 | Timing residuals for the UTMOST observations of PSR J1452–6036 between MJD 57955 and MJD 58600, using the UTMOST timing model (<i>top</i>), and a model which matches the UTMOST timing model except for an overall frequency increment of 1.1605×10^{-5} Hz (<i>middle</i>). The bottom panel shows the difference between the two sets of residuals (error bars not shown). The uncertainty on the spin frequency reported by TEMPO2 is 6×10^{-10} Hz. | 46 |
| 2.3 | Timing residuals for the synthetic dataset described in Section 2.3.1, using a timing model with (<i>top</i>) $\Delta f = 1.2106 \times 10^{-5}$ Hz which matches the injected glitch size, and (<i>bottom</i>) $\Delta f = 5 \times 10^{-7}$ Hz which is approximately 24 times smaller than the injected glitch. | 49 |
| 2.4 | Timing residuals for the synthetic dataset described in Section 2.3.1 before (<i>left panel</i>) and after (<i>middle panel</i>) adding phase jumps to account for wraps in pulse phase. The right panel shows the residuals obtained after accounting for a glitch with $\Delta f = (5.006 \pm 0.015) \times 10^{-7}$ Hz and $\Delta\phi = 0.323 \pm 0.008$ | 51 |
| 2.5 | Results of local spin frequency estimation described in Section 2.4.1. The top left panel shows the case $\tau/T \approx 2\sigma_{\text{ToA}}$, the top right panel shows the case $\tau/T \approx 10\sigma_{\text{ToA}}$, and the bottom panel shows the case where multiple observing sessions are used in each spin frequency fit. The location of the glitch is indicated by the vertical dotted line. In all cases the true size of the glitch is $\Delta f = 1.21 \times 10^{-5}$ Hz. Note the 10^2 -fold decrease in the vertical scale from the top to the bottom panels. | 54 |
| 2.6 | (<i>Top</i>) Frequency path recovered by the HMM forward-backward algorithm for the analysis of the synthetic dataset described in Section 2.3.1. The recovered glitch size is $\Delta f = 1.2078 \times 10^{-5}$ Hz. (<i>Bottom</i>) Heatmap showing the evolution of the logarithm of the posterior probability for f from the HMM analysis of the same synthetic dataset. Time increases from left to right, and at each timestep the posterior has been marginalised over \dot{f} . Values in the heatmap are clipped below -10^3 to aid readability. The favoured model, $M_2(231, 232)$, includes glitches during ToA gaps 231 and 232. | 57 |
| 2.7 | Posterior probability distributions $P(f)$ of f during the 250th ToA gap (well after the glitch at t_{231}) for twelve synthetic datasets with randomised glitch epochs in the interval $\text{MJD } 58599 < t_g < \text{MJD } 58603$ as described in Section 2.4.3. In all cases the true frequency deviation is 1.2107×10^{-5} Hz and the glitch occurs during the 231st ToA gap. | 58 |
| 2.8 | Histogram of R for all 300 pulsars in the UTMOST data release. | 60 |

| | | |
|------|--|----|
| 2.9 | Recovered frequency path (<i>top</i>) and posterior distribution of f (<i>bottom</i>) for the HMM analysis of the UTMOST observations of PSR J1709–4429, laid out as in Fig. 2.6. Before the glitch, which occurs during the 81st ToA gap, the posterior distribution of f is well-constrained, showing only a narrow band of support near $f = 0$ Hz. After the glitch, the posterior distribution of f has support in three distinct f regions, separated by $1/(1$ sidereal day). The parameters used in this analysis are reported in Table 2.2. | 61 |
| 2.10 | Local spin frequency estimation of the glitch size Δf for the glitch in PSR J1709–4429 at MJD 58178 using UTMOST observations, laid out as in Fig. 2.5. The error bars on the pre-glitch frequency estimates are too small to be seen. | 62 |
| 2.11 | Timing residuals for combined UTMOST and Parkes observations of PSR J1709–4429 for two timing models with glitch sizes $\Delta f/f = 54.6 \times 10^{-9}$ (<i>top</i>) and $\Delta f/f = 2429.7 \times 10^{-9}$ (<i>bottom</i>). In each plot the unphysical phase jump $\Delta\phi$ has also been adjusted to minimise the jump in residuals before and after the glitch. UTMOST observations are shown as squares, and the Parkes observations are shown as circles. In both plots, all other timing model parameters besides Δf and $\Delta\phi$ are identical to those published in the UTMOST data release. | 63 |
| 2.12 | Timing residuals for UTMOST observations of PSR J1452–6036. Glitches are indicated by the vertical dotted line, with sizes $\Delta f/f = 270 \times 10^{-9}$ (<i>left</i>), 2069×10^{-9} (<i>middle</i>), and 3868×10^{-9} (<i>right</i>). The post-glitch residuals appear to be identical, despite the significant differences in glitch sizes in the three timing models. Other timing model parameters match those released in the UTMOST data release. | 66 |
| 3.1 | Histogram of 90% frequentist upper limits on fractional glitch size for the 282 UTMOST pulsars analysed in this paper. | 85 |
| 3.2 | Histogram of size upper limits $\Delta f^{90\%}/f$ computed in Section 3.5.1 (orange) compared with sizes $\Delta f/f$ detected in the entire pulsar population in the Jodrell Bank Observatory catalogue (blue). | 87 |
| 3.3 | Histogram for 282 UTMOST pulsars of the ratio of 90% glitch size upper limit computed with a completely recovering glitch ($\Delta f_{Q=1}^{90\%}$) to the 90% upper limit computed with no recovery ($\Delta f_{Q=0}^{90\%}$). | 89 |
| 3.4 | Sequence of most likely frequencies $\hat{f}(t_n)$ (<i>left</i>) and heatmap of posterior frequency probability $\ln[\gamma_f(t_n)]$ (<i>right</i>) for the HMM follow-up analysis of PSR J0835–4510. Frequency is on the vertical axis in all panels, and the range of the vertical axis is the full range of the DOI. Note that the horizontal axes for the two panels are not exactly the same: the left panels have MJD on the horizontal axis, while the right panels have ToA gap index on the horizontal axis instead, for ease of plotting. The three rows correspond to sections 1, 2 and 3 from top to bottom as described in Table 3.5. | 95 |
| 3.5 | As in Figure 3.4, but for PSR J1257–1027. | 95 |
| 3.6 | As in Figure 3.4, but for PSR J1452–6036. | 95 |

| | | |
|------|---|-----|
| 3.7 | As in Figure 3.4, but for PSR J1622–4950. | 96 |
| 3.8 | As in Figure 3.4, but for PSR J1703–4851. | 96 |
| 3.9 | As in Figure 3.4, but for PSR J1709–4429. | 96 |
| 3.10 | As in Figure 3.4, but for PSR J1731–4744. | 96 |
| 3.11 | As in Figure 3.4, but for PSR J1740–3015. | 96 |
| 3.12 | Heatmap of posterior frequency derivative probability $\gamma_{\dot{f}}(t_n)$ for the HMM follow-up analysis of J1731–4744. | 97 |
| 4.1 | Illustrative heatmap showing the hidden state posterior $\gamma_{q_i}(t_n)$ for a single timestep in the analysis of the glitch event described in Section 4.3. The color scale reflects the posterior density at the 85th timestep, $\gamma_{q_i}(t_{85})$. The plotted range in f and \dot{f} reflects the complete allowed space of hidden (f, \dot{f}) states in this analysis, and the hidden state posterior is well-localised within this range. | 102 |
| 4.2 | Sequence of $\ln[K_1(k)]$ for each ToA gap in the section of data containing the glitch described in Section 4.3. A clear peak is visible at the 81st ToA gap, corresponding to MJD 48635.6 – MJD 48636.6. The threshold $\ln(K_{\text{th}}) = 1.15$ is shown as the dashed line. | 106 |
| 4.3 | (<i>Top</i>) Heatmaps of the marginalised posterior distributions of f (<i>left</i>) and \dot{f} (<i>right</i>) for the data section containing the newly discovered glitch described in Section 4.3.2. (<i>Bottom</i>) Sequence of <i>a posteriori</i> most probable f (<i>left</i>) and \dot{f} (<i>right</i>) states over time, for the same data section. The values of f and \dot{f} are relative to the pre-glitch timing solution, hence they are flat and centred on zero in the pre-glitch region. | 108 |
| 4.4 | Phase residuals from TEMPO2 before (<i>top</i>) and after (<i>bottom</i>) fitting for the new small glitch at MJD 48636 with TEMPONEST. The location of the glitch is marked by the vertical dotted line. Note that the post-fit residuals are not flat, because the TEMPONEST fit includes a timing noise component which is not subtracted from the residuals. | 109 |
| 4.5 | (<i>Top</i>) Heatmaps of the marginalised posterior distributions of f (<i>left</i>) and \dot{f} (<i>right</i>) for the data section containing the 1991 glitch described in Section 4.4.1. (<i>Bottom</i>) Sequence of <i>a posteriori</i> most probable f (<i>left</i>) and \dot{f} (<i>right</i>) states over time, for the same data section. The values of f and \dot{f} are relative to the pre-glitch timing solution, hence they are flat and centred on zero in the pre-glitch region. | 111 |
| 4.6 | (<i>Top</i>) Heatmaps of the marginalised posterior distributions of f (<i>left</i>) and \dot{f} (<i>right</i>) for the data section containing the 1999 glitch described in Section 4.4.2. (<i>Bottom</i>) Sequence of <i>a posteriori</i> most probable f (<i>left</i>) and \dot{f} (<i>right</i>) states over time, for the same data section. The values of f and \dot{f} are relative to the pre-glitch timing solution, hence they are flat and centred on zero in the pre-glitch region. | 112 |
| 4.7 | Frequentist HMM upper limits on the fractional size of undetected glitches $\Delta f^{90\%}/f$ across the 24 years of data analysed here. | 113 |

| | | |
|-----|--|-----|
| 5.1 | Sky locations of the 173 pulsars observed in the UTMOST-NS survey, shown in RA in hours and DEC in degrees. The symbol size is proportional to mean flux density. Vertical grey zones show groupings of pulsars based on proximity in RA. A single black symbol in such a zone indicates a pulsar which can be timed daily; green/blue symbols show two pulsars which are timed on alternative days; green/blue/red shows pulsars timed every third day on a three day cycle, up to a maximum of six pulsars timed over six transit days in busy regions of the Galactic plane (particularly between 15 and 20 hours RA, where the plane lies approximately north-south on the sky and there are many pulsars transiting the meridian). The horizontal width of coloured boxes containing each pulsar indicates the integration time (and is on the same scale as the RA axis, i.e. 1 hour in RA is 1 hour of elapsed time). | 121 |
| 5.2 | Before (<i>top row</i>) and after (<i>bottom row</i>) example demonstrating successful RFI cleaning using XPROF for the observation of PSR J1534–5334 taken at 2021-10-08-04:29:08, with the frequency-phase (<i>left</i>) and time-phase (<i>middle</i>) heatmaps, and the fully integrated pulse profile (<i>right</i>). Significant contamination from multiple cell network transmissions (825–830 and 840–845 MHz) is visible in the pre-cleaning plots, and largely absent in the post-cleaning plots. | 122 |
| 5.3 | Timing residuals from UTMOST-NS observations of the millisecond pulsar PSR J2241–5236. The weighted rms is $3.2 \mu\text{s}$ | 124 |
| 5.4 | Distributions of the differences in recovered timing noise parameters A_{red} and β for various combinations of methodology and dataset, as discussed in Section 5.4.2. The red stars in each panel indicate the locations of the means of the joint distributions of differences — see Table 5.3 for the numerical values of these means. When comparing the EW+NS ENTERPRISE analysis to the EW-only TEMPONEST analysis of Lower et al. (2020) (top left panel), there is no systematic shift in the recovered PSD parameters. When comparing the EW+NS ENTERPRISE analysis to EW-only ENTERPRISE analysis (top right panel), there is a tendency to recover smaller β , but no significant shift in $\log_{10} A_{\text{red}}$. The mean of the joint distribution does not differ significantly from zero when the two analyses of the EW-only data are compared (bottom panel). | 131 |
| 5.5 | Histograms comparing the expected variation in the measured value of $\ddot{\nu}$ due to timing noise, $\delta\ddot{\nu}$, against the uncertainty reported by ENTERPRISE, $\Delta\ddot{\nu}$. The results are divided into two cases depending on the estimated spectral index of the PSD, with $ \beta - 6 < 1$ in the top panel and $ \beta - 4 < 1$ in the bottom panel. Two histograms are shown in each panel, corresponding to the TN (blue) and TNLONG (orange) timing noise models. | 137 |
| 5.6 | Distribution of T_{span} required to obtain $\delta n < 1$ [(equation 5.23)], for the cases where the spectral index of the timing noise PSD satisfies $ \beta - 6 < 1$ and $ \beta - 4 < 1$. The median required T_{span} is 2.2×10^7 yr for $ \beta - 6 < 1$, much longer than the 2.8×10^2 yr for $ \beta - 4 < 1$ | 139 |

| | | |
|-----|--|-----|
| 5.7 | Illustrative example of a glitch analysis for PSR J1902+0615 using the HMM pipeline (Melatos et al., 2020), discussed in Section 5.6.1. The left panel shows the results of the model selection procedure, with log Bayes factors $K_1(k)$ (comparing one-glitch models against the no-glitch model; k indexes the ToA gap containing a glitch) and $K_2(57, k)$ (comparing two-glitch models against the best one-glitch model) plotted versus k (matched to MJD) in blue and orange respectively. The peak one-glitch log Bayes factor is 15.9, well above the threshold of 1.15, while the peak two-glitch log Bayes factor is 0.32. The middle and right panels show heatmaps of the posterior distribution of the hidden states ν and $\dot{\nu}$, with the maximum <i>a posteriori</i> tracks overlaid in red. The estimated size of the glitch based on the ν posterior distribution is $(1.1 \pm 0.1) \times 10^{-9}$ Hz. There is no discernible jump in $\dot{\nu}$. The vertical dashed lines indicate the location of the glitch. The values of γ_ν and $\gamma_{\dot{\nu}}$ are normalised so they sum to unity at each timestep. | 143 |
| 5.8 | Distribution of 90% confidence upper limits on fractional glitch size for the online (<i>top</i>) and offline (<i>bottom</i>) glitch searches described in Sections 5.6.3 and 5.6.4 (blue histograms), compared to the observed distribution of fractional glitch sizes recorded in the JBO glitch catalogue (orange histograms). All histograms are normalised so that they integrate to one. The online upper limits include one value per pulsar, 158 total. The offline upper limits include one value per time span analysed per pulsar. When a data set has been split into pre- and post-glitch sections or has been subdivided because of strong timing noise, there are multiple time spans per pulsar – in total there are 173 offline upper limits. The JBO values include all glitches with positive $\Delta\nu$, 676 in total. In the case of the online search, the mean upper limit across the sample is $\Delta\nu^{90\%}/\nu = 4.6 \times 10^{-8}$ with a maximum of 5.0×10^{-7} (PSR J1703–4851) and a minimum of 6.8×10^{-11} (PSR J1939+2134). In the case of the offline search, the mean upper limit is 5.3×10^{-9} , with a maximum of 1.3×10^{-7} (PSR J1709–4429) and a minimum of 2.4×10^{-12} (PSR J0437–4715). | 156 |
| 5.9 | Sensitivity of the UTMOST-NS array as a function of time, normalised against the average sensitivity during the observations taken between MJD 59598 and MJD 59652, indicated by the dashed vertical lines. | 159 |
| 6.1 | Illustrative plots of the transition matrix $A(q_j, q_i) = \Pr[q(t_{n+1}) = q_j \mid q(t_n) = q_i]$ according to the model described in Section 6.2.3. The top panel shows a heatmap of $A(q_j, q_i)$ where $q_i = [\Phi(t_n), f(t_n)]$ is an arbitrary initial state and $q_j = [\Phi(t_{n+1}), f(t_{n+1})]$ is the final state. The heatmap axes are $\Delta\Phi = \Phi(t_{n+1}) - \Phi(t_n)$ and $\Delta f = f(t_{n+1}) - f(t_n)$ in units of turns (cycles of 2π radians) and bins [of width $(2T_{\text{coh}})^{-1}$] respectively. The bottom panel shows slices of $A(q_j, q_i)$ as a function of $\Delta\Phi$ for $\Delta f = -1, 0, +1$ bins. Parameters: $T_{\text{coh}} = 5$ d, $\sigma = 1.76 \times 10^{-9} \text{ s}^{-3/2}$ | 170 |
| 6.2 | Detection probability P_d versus effective wave strain $h_{0,\text{eff}}$ for three choices of the HMM control parameter σ as well as a control experiment with zero phase tracking (see legend), as described in Section 6.3.4. | 177 |

| | | |
|-----|---|-----|
| 6.3 | Detection probability P_d versus effective wave strain $h_{0,\text{eff}}$ for three choices of the number of phase bins N_Φ as well as a control experiment with zero phase tracking, as described in Section 6.3.4. The three choices of N_Φ perform similarly. | 178 |
| 6.4 | Comparison of h_0^{sens} [Eq. (6.27)] with h_0^{sd} [Eq. (6.26)], with h_0^{sd} computed for three values of \dot{f} ($D = 5$ kpc in all cases). For the $\dot{f} = -5 \times 10^{-10}$ Hz s $^{-1}$ case, corresponding to the bound adopted in this paper (see Section 6.3.5), we have $h_0^{\text{sens}} < h_0^{\text{sd}}$ for $f \lesssim 0.8$ kHz. | 179 |
| 6.5 | Example mismatch profile $m(\dot{f})$ versus \dot{f} , as defined by Equation (6.28). The profile is generated from a single random realisation of both noise and signal, with all parameters other than \dot{f} fixed to their injected values. The injected \dot{f} value is denoted by \dot{f}_{inj} | 180 |
| 6.6 | Example sky resolution computations. The left and middle panels show example mismatch profiles $m(\alpha)$ and $m(\delta)$ respectively, calculated for the sky location of Terzan 6 at a frequency of 600 Hz for a single random realisation of noise and signal. The resolutions are 1.9 arcmin and 15 arcmin – note the difference in scale along the horizontal axes. The right panel shows the resolution in right ascension (blue curve) and declination (orange curve) as a function of frequency for the sky location of Terzan 6. Each point is derived by averaging the mismatch measured from 100 random realisations, and the curves are interpolated from the injection results by fitting a power law function, as described in the text. | 182 |
| 6.7 | Individual detector ASDs (top panel) and joint-detector log likelihoods (bottom panel) for the candidate group which survived the single-interferometer veto. The vertical dotted lines in the top panel indicate the extent of the Doppler modulation, and the horizontal dashed line in the bottom panel indicates \mathcal{L}_{th} . The cluster name and approximate ending frequency of the loudest candidate within the group are shown at the top of the figure. There is a clear disturbance in the H1 data within the Doppler modulation window, and no corresponding feature in the L1 data. This candidate group is vetoed. | 186 |
| 6.8 | Sensitivity estimates for the five targeted clusters. <i>Top panel</i> : characteristic wave strain at 95% detection probability, $h_{0,\text{eff}}^{95\%}$, as a function of signal frequency, f (in Hz), interpolated from the average estimated sensitivity depth defined in Eq. (6.29) over six 0.5 Hz bands. The curves for NGC 6397 and Terzan 6 are too close to be distinguished, as are the curves for NGC 6540 and NGC 6325. <i>Bottom panel</i> : minimum ellipticity $\epsilon^{95\%}$ as a function of f for the values of D listed in Table 6.1 and the fiducial value of I_{zz} in Eq. (6.30). The clusters are color-coded according to the legend. | 189 |
| 6.9 | Single-detector log likelihoods \mathcal{L}_X versus frequency for the ten candidate groups surviving the cross-cluster veto. The vertical dashed lines indicate the frequency of the loudest candidate in each group. The horizontal dashed lines indicate the loudest joint-detector log likelihood in each group, \mathcal{L}_J . Nine of the ten candidate groups satisfy $\mathcal{L}_X > \mathcal{L}_J$ and are therefore vetoed. The cluster name and approximate ending frequency of the loudest candidate in each group are recorded above each panel. | 194 |

| | | |
|------|--|-----|
| 6.10 | Single-detector log likelihoods \mathcal{L}_X versus frequency for the ten candidate groups surviving the cross-cluster veto. The vertical dashed lines indicate the frequency of the loudest candidate in each group. The horizontal dashed lines indicate the loudest joint-detector log likelihood in each group, \mathcal{L}_U . Nine of the ten candidate groups satisfy $\mathcal{L}_X > \mathcal{L}_U$ and are therefore vetoed. The cluster name and approximate ending frequency of the loudest candidate in each group are recorded above each panel. (Continued from Fig. 6.9). | 195 |
| 6.11 | Single-detector ASDs (units: $\text{Hz}^{-1/2}$) and log likelihoods versus frequency (units: Hz) for the ten candidate groups surviving the cross-cluster veto. The vertical, dotted lines in the ASD plots indicate the extent of the Doppler modulation of the candidate group. The horizontal, dashed lines in the log likelihood plots indicate the value of \mathcal{L}_{th} for the candidate's cluster. Seven of the ten candidate groups show a clear disturbance in one of the detector ASDs. The cluster name and approximate ending frequency of the loudest path in each candidate group are recorded above each panel. | 196 |
| 6.12 | Single-detector ASDs (units: $\text{Hz}^{-1/2}$) and log likelihoods versus frequency (units: Hz) for the ten candidate groups surviving the cross-cluster veto. The vertical, dotted lines in the ASD plots indicate the extent of the Doppler modulation of the candidate group. The horizontal, dashed lines in the log likelihood plots indicate the value of \mathcal{L}_{th} for the candidate's cluster. Seven of the ten candidate groups show a clear disturbance in one of the detector ASDs. The cluster name and approximate ending frequency of the loudest path in each candidate group are recorded above each panel. (Continued from Fig. 6.11) | 197 |
| 7.1 | Demonstration of the use of HMMs for discovering binary pulsars, as discussed in Section 7.2.2. The three panels show spectrograms for a loud (<i>left</i>) and quiet (<i>middle and right</i>) injection of a 173 Hz pulsar in a 80-minute circular orbit. The injected signal is clear in the loud injection, and the blue track is the recovered Viterbi path (offset artificially by 0.5 Hz to make it visually distinguishable from the track in the spectrogram). The middle panel shows the spectrogram for the quiet injection with the same orbital parameters and with no Viterbi tracks overlaid, and the right panel shows the same spectrogram along with the recovered Viterbi path in red and the true path shown as the dashed line. . . . | 207 |
| B.1 | Comparison of Viterbi scores \mathcal{S} obtained for $T_{\text{coh}} = 5 \text{ d}$ and 10 d in a set of 200 injections with \dot{f} chosen at random every 7 days in a uniform range centred on zero with extent $\pm 1.55 \times 10^{-18} \text{ Hz s}^{-2}$ (<i>top</i>), $\pm 0.76 \times 10^{-18} \text{ Hz s}^{-2}$ (<i>middle</i>), and $\pm 3.1 \times 10^{-18} \text{ Hz s}^{-2}$ (<i>bottom</i>). The dashed line indicates $\mathcal{S}_{5 \text{ d}} = \mathcal{S}_{10 \text{ d}}$ | 215 |

| | | |
|-----|---|-----|
| B.2 | Heatmaps of log likelihood \mathcal{L} for a fine-grid search in α , δ , and \dot{f} around the candidate signal in Terzan 6 at $f \approx 702$ Hz found in an early iteration of the globular cluster CW search. The top two panels present the results for $T_{\text{coh}} = 5$ d, while the bottom two present the results for $T_{\text{coh}} = 10$ d. The starred quantities \dot{f}^* , α^* , and δ^* are the values corresponding to the loudest candidate in each search. The maximum Viterbi scores in each case are given in the titles of the subfigures. The candidate is rejected because the 10 d Viterbi scores do not exceed the 5 d scores – the features of the heatmaps are shown for illustrative purposes but are not the basis for any assessment of the candidate. | 217 |
| B.3 | As in Figure B.2 but for the candidate in NGC 6325 at $f \approx 320$ Hz. | 218 |
| B.4 | As in Figure B.2 but for the candidate in NGC 6325 at $f \approx 293$ Hz. | 219 |
| B.5 | As in Figure B.2 but for the candidate in NGC 6544 at $f \approx 588$ Hz. | 220 |
| B.6 | As in Figure B.2 but for the candidate in NGC 6397 at $f \approx 491$ Hz. | 221 |
| B.7 | As in Figure B.2 but for the candidate in NGC 6540 at $f \approx 460$ Hz. | 222 |

List of tables

| | | |
|-----|---|----|
| 1.1 | Spearman correlation coefficients ρ and corresponding p -values calculated between the listed parameters and the rms timing residuals δt_{rms} for the 597 measurements reported by Keith et al. (2024). Scatter plots of the parameters against δt_{rms} are shown in Figure 1.3. | 15 |
| 1.2 | Details of a cross-section of published CW searches, covering the three categories discussed in Section 1.7.3. Here we focus on recent searches and searches directly relevant to this thesis, i.e. those using some flavour of Viterbi and those targeting particular sky regions. See Wette (2023) for a comprehensive discussion of past CW searches. Upper limits on h_0 are conditional on the details of the analysis and should not be compared directly between searches. If f_{GW} is listed as “Various” the search comprises a number of distinct narrow frequency ranges informed by electromagnetic observations of the targets’ frequencies. Abbreviations used – Gal. Cen: Galactic Center; LC: loosely coherent; SNR: supernova remnant; AMXP: accreting millisecond X-ray pulsar; NB: narrowband; NS: neutron star. References for methods: ^a Dergachev (2010); ^b Jaranowski et al. (1998); ^c Piccinni et al. (2020); ^d Suvorova et al. (2016); ^e Suvorova et al. (2017); ^f Pitkin et al. (2017); ^g Astone et al. (2012); ^h Astone et al. (2014); ⁱ Krishnan et al. (2004); ^k Bayley et al. (2019); ^l Covas and Sintès (2019); ^m Dergachev and Papa (2020). | 32 |
| 2.1 | Estimated glitch parameters for the glitch in PSR J1709–4429 at MJD 58178. The glitch model parameters are defined in equation (2.15). | 64 |
| 2.2 | HMM parameters for the analysis (described in Section 2.4.3) of the synthetic dataset described in Section 2.3.1 and the PSR J1709–4429 analysis described in Section 2.5.1. Parameters marked with an asterisk are from the UTMOST data release. | 70 |
| 3.1 | Observational statistics for the first UTMOST public data release. All quantities are calculated on a per-pulsar basis, and we take the minimum, mean and maximum over the complete set of pulsars in the data release. | 73 |
| 3.2 | Domain of interest of physical parameters and HMM control parameters. | 75 |
| 3.3 | Properties of the initial glitch candidates detected in the first UTMOST data release. | 80 |

| | | |
|-----|---|-----|
| 3.4 | Properties of the detected glitches confirmed by follow-up analyses. The fractional glitch sizes recovered in the HMM analysis are denoted by $\Delta f/f$, while the values reported by Lower et al. (2020) are denoted by $(\Delta f/f)_{\text{lit}}$. Phase ambiguity due to periodic observational scheduling prevents inferring $\Delta f/f$ for PSR J1452–6036 (see Section 3.4.2). | 81 |
| 3.5 | Data segmentation in the PSR J0835–4510 analysis. | 82 |
| 3.6 | 90% frequentist upper limits on the undetected glitch size in the seven UTMOST pulsars for which at least one glitch is detected by the HMM. . | 86 |
| 3.7 | Top and bottom five $-\dot{f}\langle\Delta t\rangle$ values among pulsars observed by UTMOST. The Pearson correlation coefficients between $\Delta f/f$ and the forward (r_+) and backward (r_-) waiting times are also listed. The objects in the top (bottom) half of the table are more likely to exhibit cross-correlations between glitch size and forward (backward) waiting time. . | 90 |
| 4.1 | Details of the data sectioning used for the HMM glitch search. Horizontal rules indicate the location of known glitches. The asterisk indicates the data section containing the newly discovered glitch described in Section 4.3. | 104 |
| 4.2 | Glitch model parameters returned by the HMM (via LMFIT) and TEMPONEST for the new small glitch described in Section 4.3.2. In both cases the glitch epoch is taken to be immediately after the last pre-glitch ToA, to match the HMM. With TEMPONEST we do not fit for t_g , instead fitting for the unphysical parameter $\Delta\phi$. The given ranges are the 90% confidence intervals in the LMFIT case, and 90% credible intervals in the TEMPONEST case. | 107 |
| 4.3 | Glitch model parameters returned by the HMM (via LMFIT) and TEMPONEST for the small glitch of 1991 described in Section 4.4.1. In the second and third columns the glitch epoch is taken to be immediately after the last pre-glitch ToA, to match the HMM. With TEMPONEST we fit for $\Delta\phi$ instead of t_g . The given ranges are the 90% confidence intervals in the second and fourth columns, and 90% credible intervals in the third column. . | 110 |
| 4.4 | Glitch model parameters returned by the HMM (via LMFIT) and TEMPONEST for the small glitch of 1999 described in Section 4.4.2. In the second and third columns the glitch epoch is taken to be immediately after the last pre-glitch ToA, to match the HMM. With TEMPONEST we fit for $\Delta\phi$ instead of t_g . The given ranges are the 90% confidence intervals in the second and fourth columns, and 90% credible intervals in the third column. . | 112 |
| 5.1 | Prior ranges used in ENTERPRISE analyses described in Section 5.3. All prior distributions are (log-)uniform. For a timing model parameter p , the best-fit value of the parameter returned by an initial TEMPO2 fit is denoted by p^* , and the associated uncertainty is denoted by Δ_p . The rms of the phase residuals from the TEMPO2 fit is denoted by σ_{rms} | 126 |

| | | |
|-----|---|-----|
| 5.2 | Estimates of the parameters determining χ_{RN} [equation (5.8)], as discussed in Section 5.4.1. In this work we consider two cases, one with $\chi_{\text{RN}} \propto f_{\text{low}}^{-\gamma}$ (second column), where γ is a free parameter to be inferred, and one with $\chi_{\text{RN}} \propto f_{\text{low}}^{-(\beta-1)/2}$ (third column). The fourth and fifth columns present the equivalent results from Shannon and Cordes (2010) and Lower et al. (2020), respectively. Error bars are 95% credible intervals for the quoted results from this work and from Lower et al. (2020), and 2- σ confidence intervals for the results of Shannon and Cordes (2010). | 129 |
| 5.3 | Mean differences in recovered timing noise parameters, $\langle \Delta \log_{10} A_{\text{red}} \rangle$ and $\langle \Delta \beta \rangle$, for three pairings of dataset and method. The dataset is either the combined UTMOST-EW and UTMOST-NS dataset (EW + NS) or the UTMOST-EW data alone (EW). The method is either ENTERPRISE (ENT) or TEMPONEST (TEM). The (EW, TEM) timing noise parameters are taken from the analysis of Lower et al. (2020). The complete distributions of timing noise parameter differences are shown in Figure 5.4. The quoted uncertainty on each mean difference is the standard error of the mean. | 132 |
| 5.4 | Pulsars with $\delta n < 1$, where δn is the expected variation due to timing noise in the measured value of n , as discussed in Section 5.5.1. The values of $\log_{10} A_{\text{red}}$, β , $\dot{\nu}$ and n are taken from the TNLONGF2 model, and the values of $\delta \dot{\nu}$ and δn are calculated according to equations (5.22) and (5.23). The quoted uncertainties on $\dot{\nu}$ and n are 95% credible intervals. | 138 |
| 5.5 | Measurements of non-zero $\dot{\nu}$ when TNLONG models are excluded from consideration, as discussed in Section 5.5.3. The measured value of $\dot{\nu}$ is given in the second column, while the implied value of n [equation (5.13)] appears in the third. The estimated properties of the timing noise PSD, A_{red} and β , are given in the fourth and fifth columns. The model chosen according to the criteria defined in Section 5.3.2 is listed in the penultimate column, and the final column gives the log Bayes factor between this chosen model and the model with the highest evidence which does not contain a $\dot{\nu}$ term. | 140 |
| 5.6 | HMM glitch detector parameters used in the online glitch detection analyses described in Section 5.6.3. Most parameters are fixed, but σ depends on the average ToA gap length in a 30-ToA window, denoted by \bar{z} | 146 |
| 5.7 | Glitch candidates identified by the offline HMM glitch search described in Section 5.6.4. The detection threshold is $\ln K_{\text{th}} = 1.15$. Note that large ($\Delta \nu \gtrsim 10^{-7}$ Hz) glitches are not included here, as the datasets for those pulsars are broken into pre- and post-glitch sections to allow for smaller DOIs and hence more sensitive searches. We indicate whether the glitch was previously identified, either as part of the UTMOST-NS online glitch detection pipeline or through by-eye inspection of the timing residuals, and whether the glitch candidate is vetoed by the procedure described in Section 5.6.4. | 148 |
| 5.8 | Timing properties of all glitches analysed in this work. Parameter estimates are from ENTERPRISE analyses, uncertainties are 95% credible intervals. The glitches are discussed in detail in Section 5.6.5. | 149 |

| | | |
|-----|--|-----|
| 6.1 | Cluster parameters for the five clusters targeted in this paper, from Harris (1996). The sky position (columns two and three) refers to the centroid of the light distribution. The core radius is denoted by r_c , the central luminosity density is denoted by ρ_c , and the distance is denoted by D . The final column is the figure of merit used to rank the clusters, as discussed in Section 6.3.1. | 174 |
| 6.2 | Injection and recovery parameters for the synthetic data tests used to motivate the choices of HMM control parameters as described in Section 6.3.4. $\mathcal{U}(a, b)$ denotes a uniform distribution between a and b | 176 |
| 6.3 | Number of sky templates and \mathcal{L}_{th} values for each target cluster as discussed in Section 6.4.1, chosen to give a false alarm rate of one per cluster on average. Targets are ordered by the number of sky templates. | 184 |
| 6.4 | Results of applying the veto procedures described in Sections 6.4.2 and 6.4.3. Each column gives the number of candidates surviving after the corresponding veto. The abbreviated column headings are: disturbed sub-band (DS), known-lines (KL), cross-cluster (CC), single-interferometer (IIFO) and unknown lines (UL). Note that the number of pre-veto candidates is a lower bound, as the candidate list from any sub-band which produces more than 1.67×10^5 candidates is not saved and is instead marked as vetoed by the disturbed sub-band veto (see Section 6.4.2). | 187 |
| 6.5 | Estimated sensitivity depths $\overline{D}^{95\%}$ for the five targeted clusters. The value of $\overline{D}^{95\%}$ is calculated by taking the mean of the $D^{95\%}$ values [see Eq. (6.29)] computed in six sub-bands. | 188 |
| B.1 | Details of the six candidates arising from an earlier iteration of the search for CW emission from globular clusters detailed in Chapter 6. For each candidate we give the corresponding cluster, the right ascension and declination (J2000), the ending frequency of the loudest Viterbi path, the value of \dot{f} for the loudest candidate, and the log likelihood \mathcal{L} | 216 |

*The big star is falling
Through the static and distance
A farewell transmission
Listen...*

Jason Molina

1

Introduction

1.1 | Neutron stars

Neutron stars are fascinating objects, home to some of the most extreme conditions in the late Universe. This is a thesis about looking at and for neutron stars, and we begin with a whirlwind tour of basic physical facts about neutron stars and their observational properties. Aspects more directly relevant to this thesis will be discussed in greater detail in later sections.

1.1.1 | Essential physics

Neutron stars are primarily formed in the supernovae of stars in a suitable mass range — those heavy enough to overcome the electron degeneracy pressure which supports white dwarfs and precipitate core-collapse, but not so massive that this core-collapse leads to the formation of a black hole.

If nothing else, they are extreme objects in a gravitational sense: neutron stars have masses of approximately one to two solar masses (M_{\odot}), and the Schwarzschild radius of a $1 M_{\odot}$ object is approximately 3 km, only a few times smaller than the typical neutron star radius. Their density is correspondingly enormous — the core densities of neutron stars exceed 10^{15} g/m^3 , beyond even the densities found in atomic nuclei. The composition of neutron stars is, in the very crudest approximation, suggested by their name: they are (mostly) made of neutrons. However, this ignores the reality of a much more complicated structure which is far from completely understood.

Neutron stars are stratified into several layers. In the outermost layer, the *crust*, the density is relatively low (though still far beyond the reach of terrestrial materials!) — this layer is composed primarily of a lattice of neutron-rich nuclei, immersed in an electron gas (Chamel and Haensel, 2008). Deeper towards the center, as the density increases beyond roughly $4 \times 10^{11} \text{ g cm}^{-3}$, it is energetically favorable for a portion of the neutrons to “drip” out of the nuclei (Rüster et al., 2006), and this inner crust is thus augmented

by a neutron gas which is expected to be superfluid (Chamel, 2017). The presence of superfluid components in the neutron star interior plays a significant role in some of the phenomena observed in pulsar timing which are of particular interest in this thesis, as we shall see in Sections 1.4 and 1.5. The transition from the crust to the core occurs at densities comparable to the nuclear saturation density, $\rho_0 \approx 2.8 \times 10^{14} \text{ g cm}^{-3}$, at which point the nuclei, which have already begun to deform in the deepest parts of the inner crust into a variety of exotic shapes known as “nuclear pasta” (Blaschke and Chamel, 2018), completely dissolve. In this outer core region, the majority of the matter is in the form of a superfluid neutron component, with a small fraction (roughly a few percent) in the form of a superconducting proton state (Link, 2003; Alford and Good, 2008; Wood and Graber, 2022). In the outer core, the proton superconductor is typically believed to be type-II, so that the magnetic flux is quantized into an array of microscopic “flux tubes” (Baym et al., 1969a). The interactions of these flux tubes with the quantised vortices of rotation which form in the neutron superfluid component are expected to be strong, and the dynamics involved may have a significant effect on the rotational evolution of the star (Bhattacharya and Srinivasan, 1991; Ruderman et al., 1998; Drummond and Melatos, 2018; Thong et al., 2023). In the inner core, where the density significantly exceeds ρ_0 , the composition is highly uncertain. It is likely that there is a transition from the hadronic regime, where quarks are confined in neutrons, protons, etc., to the quark matter regime, where the hadrons melt into a sea of strongly interacting u , d , and s quarks (Baym et al., 2018).

The magnetic fields of neutron stars are also extreme, with typical strengths between 10^8 G and 10^{15} G (Manchester et al., 2005). The mechanism by which these fields are formed and maintained is not well understood, but simple flux conservation arguments are likely insufficient to explain the full range of observed magnetic field strengths, and some internal dynamo process is needed to amplify the fields (Thompson and Duncan, 1993; Ferrario et al., 2015). The simplest picture of a magnetised neutron star is as a magnetic dipole rotating in vacuum – but reality is significantly more complex. For one thing, the neutron star is not rotating in vacuum – in the classical picture, the rotation of the magnetic field accelerates charges off the surface of the neutron star, generating a plasma surrounding the star: the *magnetosphere* (Goldreich and Julian, 1969). The radio emission by which pulsars were first discovered and by which the majority of known neutron stars today are observed (see Section 1.1.2) is likely to originate from processes in this region, but the details are not well understood (Philippov and Kramer, 2022). The characteristics of the radio emission thus serve as a valuable probe into the nature of the magnetosphere and the structure of the magnetic field.

1.1.2 | Observational properties

The vast majority of known neutron stars are observed as *pulsars* – astronomical objects which emit a steady sequence of pulses in some region of the electromagnetic spectrum. The first pulsar was discovered in 1967 by Jocelyn Bell Burnell in the course of a radio survey aimed at measuring interplanetary scintillation of compact radio sources (Hewish et al., 1968), and although the origin of the signal was initially unclear, it was quickly understood that a rotating neutron star was the only plausible source, with the pulsations generated by a “lighthouse effect” due to the misalignment of the magnetic axis (with

which the radiation beam is aligned, see Section 1.1.1) and the star’s rotation axis (Pacini, 1967; Gold, 1968). Since this initial discovery, many more pulsars have been discovered – at the time of writing, the ATNF Pulsar Catalogue¹ (Manchester et al., 2005) lists 3473 known pulsars, and new pulsars continue to be discovered both in ongoing surveys and in analyses of legacy data. Although originally discovered in the radio, pulsars are found across essentially the full electromagnetic spectrum, from radio observations at 10 MHz (Stappers et al., 2011) up to gamma rays with energies of 20 TeV (H.E.S.S. Collaboration et al., 2023). Not all known neutron stars are observed via pulsed emission. Some are observed as components of interacting binary systems, where the X-ray emission due to accretion can be observed, but pulsations are not necessarily detected (Sazonov et al., 2020; Fortin et al., 2023). Additionally, a select few nearby neutron stars are seen in X-rays via their thermal emission – the so-called “Magnificent Seven” (Haberl, 2007). However, the observational focus of this thesis is on pulsars, and we will not have more to say about these non-pulsar neutron stars here.

Pulsars can broadly be divided into three categories, depending on the source of the pulsed emission: *rotation-powered*, *accretion-powered*, and *magnetars*.

- *Rotation-powered pulsars* comprise the majority of known pulsars, and their emission is powered by the rotational energy of the neutron star. They are primarily observed in radio, but examples of higher-energy emission extending all the way up to gamma rays are known. They are further subdivided into two subcategories. *Canonical* pulsars are those which have not experienced significant interaction with other objects since their birth, and typically have periods between 0.1 s and 10 s. By contrast, *recycled* pulsars are old pulsars which have spent time accreting sufficient matter from a binary companion to be spun up to short periods (Alpar et al., 1982; Archibald et al., 2009), usually less than ~ 10 ms, hence they are also known as *millisecond* pulsars (MSPs). Note however that when they are observed as recycled pulsars they are no longer accreting (except for a subclass of transitional objects which are observed to switch between accreting and non-accreting states, e.g. Papitto et al., 2013). A plot showing the distribution of rotation-powered pulsars in terms of their period P and period derivative \dot{P} is shown in Figure 1.1.
- The emission of *accretion-powered* pulsars is, as the name suggests, derived from the energy of accreted matter falling onto the neutron star surface, which is funneled by the magnetic field onto the poles and forms “hotspots” which emit strongly in X-rays. These pulsars display a wide variation in behaviour depending on the mass of the companion. Low-mass X-ray binaries (LMXBs) are formed when a companion star of mass $\lesssim 1M_{\odot}$ overflows its Roche lobe, and an accretion disk is formed around the neutron star. These systems are the precursors to MSPs, and in systems where a pulsation frequency can be measured it is often found to be increasing with time due to the accretion torque, i.e. these systems are actively being recycled (e.g. Riggio et al., 2008; Riggio et al., 2011; Bult et al., 2021). By contrast, high-mass X-ray binaries (HMXBs) contain a more massive companion and the accretion in this case is primarily driven by stellar winds, rather than Roche

¹atnf.csiro.au/research/pulsar/psrcat/

lobe overflow (Reig, 2011). The pulsation periods in this case are much slower, ranging from 0.1 s to 10^4 s (Fortin et al., 2023).

- *Magnetars* are powered by the decay of their magnetic fields, which are the strongest known among the pulsar population – up to 10^{15} G (Olausen and Kaspi, 2014). Though primarily observed as X-ray or gamma-ray sources, a handful are also observed to emit in the radio (Camilo et al., 2006; Kaspi and Beloborodov, 2017). They are typically slow rotators, with periods between 1–10 s. Magnetar emission is extremely dynamic – they exhibit short (millisecond-duration) bursts in both their high-energy emission (Kouveliotou et al., 1999; van der Horst et al., 2012) and in the radio (Bochenek et al., 2020; CHIME/FRB Collaboration et al., 2020), as well as longer-lived flares and outburst episodes (Hurley et al., 2005; Lower et al., 2023b).

In this thesis the focus is on rotation-powered pulsars, for which the primary mode of observation is *pulsar timing*, which gives an extremely precise method of tracking the rotational behaviour of a pulsar over timescales of many years. Radio pulsar timing is discussed in further detail in Section 1.2.

1.1.3 | Neutron stars in globular clusters

Although the majority of known pulsars formed from stars residing in the Galactic disk, there is an interesting sub-population of pulsars which formed not in the disk but in globular clusters (GCs). GCs are dense star clusters, with core luminosity densities exceeding 10^5 solar luminosities per cubic parsec and hosting extremely old stellar populations, with estimated ages on the order of 10 Gyr (Baumgardt, 2017). There are approximately 150 known GCs associated with the Milky Way (Harris, 1996). While much of this thesis deals with canonical pulsars, which are rarely found in globular clusters (as discussed below), Chapter 6 presents a search for gravitational radiation from fast-spinning neutron stars in five nearby GCs, and hence we devote this section to a brief discussion of the observed population of GC pulsars.

Globular clusters have long been recognised as hosting many more LMXBs per unit mass than the Galactic disk (Katz, 1975; Clark, 1975), and given that LMXBs produce MSPs (see Section 1.1.2), one might well expect that GCs host an excess of MSPs compared to the field, and indeed this is confirmed by observation (Camilo and Rasio, 2005; Ransom, 2008; Ye et al., 2019). At the time of writing there are 330 known pulsars among 44 globular clusters² – of these 330 pulsars, 275 (83%) have spin periods less than 10 ms. This is a much higher proportion than the overall pulsar population, for which only 14% of pulsars have $P < 10$ ms.

The preponderance of MSPs in globular clusters is due not only to the high rate of LMXB formation, but the age of their stellar populations. Stars massive enough to undergo core collapse and form neutron stars have lifetimes of only ~ 10 Myr (Crowther, 2012), and given the age of GC populations these “first-generation” neutron stars are also extremely old. However, the spin-down timescale of the canonical pulsars is on the order of 10^7 yr, and hence these non-recycled (i.e. slow) pulsars have long since spun

²www3.mpifr-bonn.mpg.de/staff/pfreire/GCpsr.html




Figure 1.1: P - \dot{P} diagram of the full pulsar population generated from the ATNF pulsar catalogue (Manchester et al., 2005) using `PSRQPY` (Pitkin, 2018). Pulsars in binaries are circled in red. Lines of constant characteristic age ($\tau_c = P/2\dot{P}$) and surface magnetic field [$B_{\text{surf}} = 3.2 \times 10^{19} \text{ G} (P/1 \text{ s})^{1/2} \dot{P}^{1/2}$] are shown as dashed and dotted lines respectively. The shaded region in the bottom right is below the “death line” as defined in Equation 4 of Zhang et al. (2000), approximately where pulsar radio emission is expected to cease.

down to the point where their radio emission turns off, and are no longer visible. However, although MSPs dominate the population of GC pulsars, slow, apparently young pulsars with periods greater than 0.1 s are also observed (Lyne et al., 1993; Biggs et al., 1994; Boyles et al., 2011). There are several proposed formation channels for these slow pulsars, including electron capture supernovae of white dwarves (Ivanova et al., 2008; Boyles et al., 2011; Tauris et al., 2013; Kremer et al., 2023) and disruption of the recycling process by a close encounter with another object (Verbunt and Freire, 2014).

Not only is the overall makeup of the pulsar population in GCs markedly different from the distribution observed in the Galactic field, the properties of GC MSPs in particular are distinct from the field population (Lee et al., 2023). Most strikingly, GCs host a higher proportion of isolated MSPs than the field – 42% of GC MSPs are isolated, compared to 20% of field MSPs³. Verbunt and Freire (2014) argue that the production of isolated MSPs in globular clusters is due to the high rate of encounters experienced by binary systems in GCs over their lifetime. Moreover, they find evidence for variation in the production rate of isolated MSPs between globular clusters, and present a characteristic figure of merit which tracks the rate at which binary systems experience further encounters and is based on the structural parameters of the cluster,

$$\gamma \propto \frac{\sqrt{\rho_c}}{r_c}, \quad (1.1)$$

where ρ_c is the core luminosity density and r_c is the core radius. Using the population statistics known at the time, they showed that clusters with high γ tend to produce a higher proportion of isolated MSPs than those with low γ , and subsequent discoveries of new GC pulsars appear to bear this out (e.g. Ridolfi et al., 2022; Vleeschower et al., 2022; Yin et al., 2024). This relation between γ and isolated MSPs in the population is not exact – e.g. see Chen et al. (2023) for an discussion of the MSP population in the massive globular cluster ω -Cen, which has a low γ but appears to have an MSP population dominated by isolated systems. The choice of target clusters in the CW search presented in Chapter 6 is heavily informed by these considerations; see Section 6.3.1.

1.2 | Radio pulsar timing

As discussed in Section 1.1.2, almost all neutron stars are observed as pulsars, producing electromagnetic signals with astonishingly precise periodicity, typically assumed to be sourced by radiation generated above the magnetic poles of the neutron star, and so the pulse sequence is directly tied to the rotation of the star (the magnetic field itself being locked to the crust of the star). The extreme precision of the observed periodicity motivates the following question: can we build a useful model which predicts the arrival of each pulse from a given pulsar at a telescope on Earth? This is the essential game of *pulsar timing*, and the answer turns out to be yes. In this section we will go through how this is achieved in some detail, focusing on the timing of radio pulsars⁴.

³www.astro.umd.edu/~eferrara/pulsars/GalacticMSPs.txt

⁴While pulsar timing is also routinely performed at other wavelengths, especially in X-rays and gamma rays, the details around how the data are acquired and processed are beyond the scope of this thesis.

1.2.1 | Data acquisition

For the purposes of this thesis, the starting point of an observation of a radio pulsar is a digitised time series $V(t)$ recording the voltage⁵ received by the telescope from a certain direction on the sky within a certain radio frequency (RF) band (determined by the telescope configuration). This already elides many details! The formation of this time series is not trivial, and happens differently depending on the configuration of the telescope (Lower, 2022). Although we are considering only one polarisation, frequently two orthogonal polarisation modes are recorded – information about the polarisation state of the emission is extremely valuable for studies of pulsar emission physics (Manchester, 1971; Ruderman and Sutherland, 1975; McKinnon and Stinebring, 2000; Oswald et al., 2020), but in this thesis we will not make use of it. The RF band covered by $V(t)$ is typically centred somewhere in the range of $O(100 \text{ MHz})$ to $O(1 \text{ GHz})$, and bandwidths range from a few MHz to a few GHz.

This voltage time series is necessarily sampled at a very high rate⁶ – by the Nyquist sampling theorem, a signal of bandwidth f must be sampled at a rate exceeding $2f$. Except in rare circumstances this raw data stream is not stored. Instead, $V(t)$ is processed on the fly to produce data products which are easier to handle. In modern systems with wide bandwidths, the first step is to pass $V(t)$ through an array of bandpass filters (a “filterbank”) to produce a new set of voltage time series, each of which cover only a fraction of the total observing bandwidth. This allows the downstream processing to be parallelised over separate compute nodes, rather than requiring ingestion of the full-bandwidth signal (e.g. Hobbs et al., 2020; CHIME/Pulsar Collaboration et al., 2021; Mandlik et al., 2024).

The most important final data product in our case is the *folded pulse profile*. If we are observing a known pulsar whose pulse period P is already measured with reasonable precision, we can break up the data stream into chunks of length P and add the intensity timeseries in each chunk together, to produce a single *integrated* pulse with much higher signal-to-noise ratio than any individual pulse in the data. In forming this integrated pulse we must contend with *dispersion* – the frequency-dependent delay in the arrival time of radio photons due to the refractive properties of the interstellar medium, in particular the presence of cold plasma along the line of sight. The dispersive delay at an observing frequency ν_{obs} relative to an undispersed pulse is given by (Kulkarni, 2020)

$$\tau_{\text{DM}}(\nu_{\text{obs}}) = \frac{a}{\nu_{\text{obs}}^2} \int_0^L n_e dl \quad (1.2)$$

where a is a constant equal to approximately $4.15 \text{ GHz}^2 \text{ cm}^3 \text{ pc}^{-1} \text{ ms}$ (Lorimer and Kramer, 2004), n_e is the electron density, and the integral is taken along the line of sight to the pulsar, which is a distance L away. The dispersion measure $\text{DM} = \int_0^L n_e dl$ is usually quoted in units of pc cm^{-3} . For an observation taken over an observing band $[\nu_1, \nu_2]$ the

⁵Note that “voltage” in this context does not refer to the electric potential difference, but an electrical signal proportional to a component of the electric field (determined by the antenna polarisation) at the receiver.

⁶Although not necessarily at a high bit depth. In some cases each sample may be digitised at only 1- or 2-bit resolution (Sarkissian et al., 2011).

degree of dispersive smearing is given by

$$\tau_{\text{DM}} = a \times \text{DM} \times \left(\frac{1}{\nu_1^2} - \frac{1}{\nu_2^2} \right) \sim 83 \text{ ms} \left(\frac{\text{DM}}{100 \text{ pc cm}^{-3}} \right) \left(\frac{\nu_1}{1 \text{ GHz}} \right)^{-2} \left(\frac{\Delta\nu/\nu_1}{10\%} \right), \quad (1.3)$$

where the final expression holds to an order of magnitude as long the fractional bandwidth $\Delta\nu/\nu_1 = (\nu_2 - \nu_1)/\nu_1$ is not too large. The dispersive smearing timescale can be a significant fraction of a pulse period, or even more than one pulse period, and thus it is essential to correct for this effect. It is possible to do this by operating on the voltages directly (if the voltages have been channelised by a filterbank this step can be done separately on each channel), convolving $V(t)$ with the inverse frequency response of the ISM – this is known as *coherent dedispersion* (Hankins and Rickett, 1975) and although it once required specialised hardware to perform, today it can be performed in software with the aid of graphics processing units (GPUs) (Ford et al., 2010; De and Gupta, 2016). If coherent dedispersion is not available, the effect of dispersion can still be corrected for incoherently by taking the channelised voltage time series and “detecting” them (converting the measured voltages to intensities), and offsetting these intensity data streams in time from one another according to equation (1.3) (Large and Vaughan, 1971; Lyne and Graham-Smith, 2012). Incoherent dedispersion suffers from a degree of dispersive smearing *within* each channel, but it is a computationally cheap alternative to the coherent approach.

Each channel in the dedispersed filterbank is folded at the pulse period to produce a time-integrated profile for each channel. This process is typically only partially performed before the result is written to storage – e.g. in a 10-minute observation, each 20-second span of data might be folded, but these 20-second foldings are not subsequently summed together before being saved. These partial foldings are referred to as “subintegrations”.

Often, however, it is useful to fully integrate in both frequency and time – i.e. sum the intensities in all subintegrations and across all frequency channels. This produces a pulse profile which is an intensity timeseries covering a single pulse period. It is a remarkable fact of pulsar timing that these fully integrated pulse profiles are extremely stable on timescales of days to years, despite the fact that individual pulses from a pulsar may vary wildly in shape and amplitude (Drake and Craft, 1968; Helfand et al., 1975; Osłowski et al., 2014; Parthasarathy et al., 2021).

We can take advantage of this profile stability to *further* reduce the data. If the integrated profile is time-tagged to high accuracy, we can ask: “when did this profile arrive at our telescope?” More specifically we can pick some reference point within the observed profile (the point of peak intensity, for example) and ask, for a given profile, what time at the telescope this reference point corresponds to for a given observation. The answer to this question can be computed in a number of ways (Wang et al., 2024), but in any case the result is a pulse *time of arrival* (ToA) which by assumption corresponds to a fixed point in the rotational phase of the pulsar (when the pulsar is “pointed towards us”, loosely speaking). It is worth emphasising that the result is *a* time of arrival, singular. While we do have the freedom to break up the data at an earlier stage in various ways, e.g. into separate frequency bands, or into a sequence of consecutive time chunks, going through the process of dedispersing, channelising, folding, integrating, and comparing the data against a reference profile will produce two things: a time of arrival, and

an uncertainty on the time of arrival. An entire observation of a pulsar is thus reduced to two numbers. The ToA is typically quoted as a Modified Julian Date (MJD), which measures time by the number of elapsed days since 0:00 November 17, 1858 AD. And yet even after this extreme compression, pulsar datasets are exquisitely information-rich by virtue of their accuracy – depending on the pulsar and the telescope, uncertainties on times of arrival range from milliseconds to hundreds of nanoseconds.

1.2.2 | Pulsar timing models

As discussed, pulsar timing datasets are made up of *times of arrival*, which are time-stamps at which pulses from a pulsar arrived at some telescope on Earth. But it is a long journey from there to here, and constructing a model which can accurately predict when pulses will arrive is an involved process – here we give only a flavour of what is involved.

Fortunately, the problem of predicting pulse arrival times in a reference frame which is at rest with respect to the pulsar (and free of any intervening material) is very easy to get *mostly* right. We make the assumption that pulse times of arrival correspond to a fixed value of the pulsar’s rotational phase ϕ – say, $\phi = 0$ – and hence aim to write down a model for the phase evolution of the pulsar, so that the times of arrival are exactly the locations of the zero crossings of ϕ . To a very good approximation, the rotational phase of a pulsar can be written as a Taylor series in the frame at rest with respect to the pulsar, with time coordinate t_p :

$$\phi(t_p) = \phi_0 + \sum_{n=1}^N \frac{1}{n!} \frac{d^n \phi}{dt^n} (t_p - t_0)^n, \quad (1.4)$$

where an expansion to second order [i.e. $\phi(t_p) = \phi_0 + \nu(t_p - t_0) + \dot{\nu}(t_p - t_0)^2/2$ where ν is the pulse frequency and $\dot{\nu}$ is its first time derivative] typically suffices to capture almost all of the phase evolution.

But we do not have the luxury of observing in the pulsar frame. Emission from a pulsar must travel across the interstellar medium, through the solar system, and into a telescope which is attached to the Earth which is both rotating upon its own spin axis and orbiting the sun. If the pulsar is in a binary with a companion object the story is complicated further, as the pulsar’s binary motion and the gravitational influence of its companion induce further corrections.

All of these considerations essentially amount to the need for a transformation between the pulsar time t_p and the time at the telescope, t_{obs} :

$$t_p = t_{\text{obs}} - \text{correction terms}, \quad (1.5)$$

where the correction terms depend on parameters such as the pulsar position on the sky, the dispersion measure, and binary orbital elements – see Edwards et al. (2006) for a detailed discussion of these effects.

Although this basic model typically describes the data to a good approximation, the difference between the predicted times of arrival and those observed (the “timing residuals”) can exhibit significant structure beyond the noise due to measurement uncertainty.

The signals received from pulsars contain information about a wide variety of phenomena which are due both to processes directly related to the pulsar and its immediate environment, and external processes which leave an imprint on the signal as it propagates to telescopes on Earth. The external processes are largely beyond the scope of this thesis – these include effects such as interstellar scintillation, which probes the spatial and temporal variation in the turbulent interstellar medium (Lyne, 1984; Stinebring et al., 2001; Reardon, 2018), and the minute changes in pulse arrival times due to the influence of low-frequency gravitational waves with frequencies of order 1 yr^{-1} (Sazhin, 1978; Detweiler, 1979; Reardon et al., 2023; Agazie et al., 2023; EPTA Collaboration et al., 2023; Xu et al., 2023). Also beyond the scope of this thesis are intrinsic effects related to the pulsar emission itself – for example, the stability of the integrated pulse profile in the majority of pulsars belies a highly dynamical process which becomes apparent in careful analysis of high time resolution observations allowing for studies of the individual pulses from subsequent rotations of bright pulsars (Drake and Craft, 1968; Backer, 1970; Knight et al., 2006; Osłowski et al., 2014; Parthasarathy et al., 2021).

The phenomena that *are* of particular interest for this thesis, and to which the following three sections are dedicated, are those related to the intrinsic rotational behaviour of the pulsar. Despite their reputation as highly precise clocks, pulsars nonetheless exhibit non-trivial rotational behaviours, which appear as signatures in the sequence of received times of arrival and provide a valuable window into neutron star physics.

1.3 | Pulsar spin down

To a very crude approximation, pulsars are rotating magnetic dipoles. Such a rotating dipole will necessarily emit electromagnetic radiation, extracting energy and angular momentum and causing the rotation of the dipole to slow down over time. The spin-down rate due to dipole radiation is given by (Lorimer and Kramer, 2004)

$$\dot{\nu} = -(9.8 \times 10^{-13} \text{ Hz s}^{-1}) \left(\frac{B}{10^{12} \text{ G}} \right) \left(\frac{\nu}{10 \text{ Hz}} \right)^3, \quad (1.6)$$

where B is the magnitude of the magnetic field at the pole of the neutron star. The values of $\dot{\nu}$ measured in the pulsar population range between $-1.3 \times 10^{-18} \text{ Hz s}^{-1}$ and $-3.8 \times 10^{-10} \text{ Hz s}^{-1}$,⁷ and dipole radiation is generally expected to be the dominant effect determining the pulsar’s frequency evolution. The dipole radiation is emitted at the rotation frequency, $\nu \sim 0.1\text{--}100 \text{ Hz}$ – the energy of this low-frequency radiation goes into “heating” the surrounding medium, and can not propagate to Earth. In some cases this heating can be observed as a pulsar wind nebula: charged particles surrounding the pulsar absorb a fraction of the deposited energy and re-radiate it across the electromagnetic spectrum (Gaensler and Slane, 2006; Lyne and Graham-Smith, 2012). Note that as the neutron star is not truly a dipole rotating *in vacuo*, the true braking behaviour is modified by the effects of particle outflows, and this tends to soften the scaling of $\dot{\nu}$ with ν (Melatos, 1997; Contopoulos and Spitkovsky, 2006; Bucciantini et al., 2006).

Magnetic dipole radiation is not the only mechanism by which pulsars may spin down. If the pulsar is emitting gravitational waves (GWs), this too will act to spin down the star

⁷We exclude *positive* measured values of $\dot{\nu}$, which are due to contamination from other effects such as accretion or motion in a gravitational potential.

over time. Gravitational waves are discussed more detail in Sections 1.6 and 1.7; for now we briefly note two particular GW emission mechanisms which may be relevant. One is mass quadrupole radiation, due to the pulsar possessing a mass distribution which is not axisymmetric about the rotation axis. In this case the spindown is described by (Riles, 2023)

$$\dot{\nu} = -(1.7 \times 10^{-19} \text{ Hz s}^{-1}) \left(\frac{\epsilon}{10^{-6}} \right)^2 \left(\frac{\nu}{10 \text{ Hz}} \right)^5, \quad (1.7)$$

where ϵ is the ellipticity, which measures the extent of the non-axisymmetry (see Section 1.7.1). The other case is emission via r -modes — these are certain oscillation modes in the neutron star interior which are unstable to GW radiation and may be driven to significant amplitudes (Andersson, 1998; Owen et al., 1998). The frequency evolution in this case is described by (Riles, 2023)

$$\dot{\nu} = -(7.70 \times 10^{-21} \text{ Hz s}^{-1}) \left(\frac{\alpha}{10^{-3}} \right) \left(\frac{\nu}{10 \text{ Hz}} \right)^7, \quad (1.8)$$

where α is the dimensionless r -mode amplitude.

The spindown evolution in all of the cases mentioned here — magnetic dipole radiation, mass quadrupole radiation, and r -mode radiation — follow the same basic form:

$$\dot{\nu} = -K\nu^n. \quad (1.9)$$

The n in equation 1.9 is referred to as the *braking index*. Magnetic dipole radiation has $n = 3$, mass quadrupole $n = 5$, and r -modes have $n = 7$. The measurement of n in a given pulsar can thus be extremely valuable in determining the dominant spindown mechanism. If the constant of proportionality K is constant, n can be measured in principle via the second frequency derivative, i.e.

$$n = \frac{\nu \ddot{\nu}}{\dot{\nu}^2}. \quad (1.10)$$

Measurement of $\ddot{\nu}$ is difficult in most cases. It has long been understood that although a straightforward least-squares fit for a $\ddot{\nu}$ term may often return a well-constrained value, these values are almost always hopelessly contaminated by long-term wandering in the phase residuals (timing noise; discussed in Section 1.4) and have no relation to the long-term spin-down behaviour of the pulsar (Hobbs et al., 2004; Chukwude and Chidi Odo, 2016). It is sometimes claimed that more sophisticated techniques which simultaneously model both the secular phase evolution and timing noise (van Haasteren et al., 2009; Lentati et al., 2014) offer a means of measuring $\ddot{\nu}$ which is robust against the presence of timing noise (Lower et al., 2020; Parthasarathy et al., 2020). The values of n measured in this way are sometimes orders of magnitude from the canonical value of 3, and may be either positive or negative. However, recent work by Vargas and Melatos (2023a) and Keith and Nițu (2023) has demonstrated that these values must also be interpreted with caution, as the timing noise model may fail to adequately capture the full extent of the effect of the timing noise on long timescales which are most relevant to the measurement of $\ddot{\nu}$. This is also explored in Section 5.5 of this thesis.




Figure 1.2: Timing residuals for the pulsar PSR J1224–6407, using observations from the UTMOST pulsar timing programme (see Jankowski et al. 2019 and Chapter 5 of this thesis). Significant long-term wandering behaviour (timing noise) is evident.

1.4 | Timing noise

As discussed in Section 1.2.2, once the pulse times of arrival have been transformed into the reference frame of the pulsar itself, the typical timing model is simple, expressed as a Taylor expansion about some reference time t_0 and frequently no more complex than

$$\phi(t_p) = \phi_0 + \nu(t_p - t_0) + \dot{\nu}(t_p - t_0)^2/2. \quad (1.11)$$

Does such a model completely describe the behaviour of pulsars, to the limits of instrumental precision? Almost always the answer is no. After fitting a model like equation (1.11), the timing residuals — the difference between the predicted and measured times of arrival — show time-correlated structure with timescales ranging from days to years, generically referred to as *timing noise*. An example is shown in Figure 1.2. Timing noise has been observed since the early days of pulsar timing (Nelson et al., 1970; Boynton et al., 1972), and a wide variety of both phenomenological and physical models have been proposed to explain this behaviour.

1.4.1 | Measuring and characterising timing noise

A common early phenomenological description was introduced by Boynton et al. (1972) which described the timing noise in terms of a random walk, which may be present in the spin-down ($\dot{\nu}$), frequency, or phase. This was successfully applied to a number of pulsars (Manchester and Taylor, 1974; Groth, 1975; Cordes and Helfand, 1980), though in other cases a random walk in a given parameter was found to be insufficient to explain the observed behaviour (Gullahorn and Rankin, 1982; Cordes and Downs, 1985; D’Alessandro et al., 1995).

Today, the dominant mode of timing noise analysis is based on an analysis of the timing noise as a Gaussian process in the timing residuals (van Haasteren et al., 2009; van Haasteren and Levin, 2013; Lentati et al., 2013). A more detailed exposition of the framework underpinning this mode of analysis is given in Appendix A, here we cover only the essential idea. The contribution of the timing noise process to the observed ToAs is modelled as a sum of sinusoids with frequencies f_j , $j = 1, 2, \dots, N$ where N is typically of order 10. Most often the frequencies take the values $f_j = j/T$ where T is the total timespan of the dataset, but other choices are possible. The power spectral density (PSD) of the timing noise contribution is usually assumed to have a power law form,

$$S(f) = \frac{A_{\text{red}}^2}{12\pi^2} \left(\frac{f}{1 \text{ yr}^{-1}} \right)^{-\beta} (1 \text{ yr}^{-1})^3. \quad (1.12)$$

The parameters A_{red} and β are the ultimate inference targets in most timing noise analyses. The power of this approach is that the timing noise contribution is estimated simultaneously with the parameters controlling the deterministic part of the phase evolution – e.g. ν , $\dot{\nu}$, sky position, binary orbital elements, etc. In principle this allows for a better understanding of the covariance between the influence of timing noise and changes in the model parameters, and hence posteriors which reflect the uncertainties introduced by the presence of timing noise. This is the origin of the claims discussed briefly in Section 1.3, that modern pulsar timing techniques are able to infer $\dot{\nu}$ in a way which is robust against timing noise contamination.

As a complementary approach to estimating the power spectrum of the timing noise – thereby marginalising over the details of the specific realisation which is observed – one can also explicitly track the evolution of the pulsar’s rotation over time. This is the approach taken in the hidden Markov model-based pulsar timing framework which is used throughout this thesis and discussed in more detail in Section 1.10. There, the pulsar frequency and frequency derivative are tracked over time, and maximum *a posteriori* estimates of their value at any given timestep can be extracted. However, in the applications of this framework so far, which largely deal with the detection of glitches, timing noise is essentially a nuisance to be dealt with, rather than the phenomenon of primary interest, and so little work has been done on extracting information about the timing noise itself from these analyses. By contrast, another state-space based approach introduced by Meyers et al. (2021a) and refined and applied to real data by Meyers et al. (2021b) and O’Neill et al. (2024) respectively, puts the nature of the timing noise at the centre of the analysis. Here a simple two-component model of the neutron star is assumed (Baym et al., 1969b), which divides the star into a superfluid interior which rotates with angular velocity Ω_s and a rigid crust which rotates with angular velocity Ω_c . A Kalman filter (Kalman, 1960) is used to track the evolution of the pair (Ω_s, Ω_c) , and also to provide a likelihood with which inference can be performed on the parameters of the model, e.g. the coupling timescale between the two components. O’Neill et al. (2024) demonstrated the application of this technique to radio observations of PSR J1359–6038 (which provide only measurements of Ω_c), and were able to constrain e.g. the coupling timescale between crust and interior to $10^{7.1+0.8}_{-0.5}$ s, and show that a two-component model is preferred over a one-component model where the whole star rotates rigidly.

1.4.2 | Timing noise across the population

Given the ubiquity of timing noise in the pulsar population, it is natural to ask how features of the timing noise correlate with other pulsar properties. The most straightforward phrasing of this question is to ask whether the “amount” of timing noise scales with the spin parameters of the pulsar, and this question has been investigated by several authors. Cordes and Helfand (1980) measured the “timing activity” of a pulsar by the rms phase residuals remaining after a second-order polynomial fit (i.e. a fit for ϕ_0 , ν , and $\dot{\nu}$). Based on a sample of 50 pulsars, they found that the timing activity is positively correlated with \dot{P} and at most weakly correlated with P . Similarly, Arzoumanian et al. (1994) parametrised the amount of timing noise by

$$\Delta_8 = \log_{10} \left(\frac{1}{6\nu} |\dot{\nu}| t^3 \right), \quad (1.13)$$

where the values of ν and $\dot{\nu}$ are measured over a timespan of $t = 10^8$ s (arbitrarily chosen). With a sample of 96 pulsars they fit (by eye) a scaling relation $\Delta_8 = 6.6 + 0.6 \log \dot{P}$ (Hobbs et al. 2010 later gave a revised scaling of $\Delta_8 = 5.1 + 0.5 \log \dot{P}$ based on a linear least-squares fit with a sample of 366 pulsars). Matsakis et al. (1997) introduced another parametrisation of the timing stability of a pulsar modelled on the Allan variance used to characterise terrestrial clocks (Allan, 1987), and defined as

$$\sigma_z(\tau) = \frac{\tau^2}{2\sqrt{5}} \langle c_3^2 \rangle^{1/2}, \quad (1.14)$$

where c_3 is the cubic coefficient in a polynomial fit to the data (i.e. the $\dot{\nu}$ of the Taylor expansion of ϕ) and the angle brackets denote averaging over subsequences of length τ . Hobbs et al. (2010) showed that $\sigma_z(10 \text{ yr})$ is weakly correlated with ν and strongly correlated with $\dot{\nu}$, essentially confirming the conclusions of Cordes and Helfand (1980). Similarly, Urama et al. (2006) took measurements of $\dot{\nu}$ to be a direct indicator of timing noise and demonstrated a strong correlation between $\dot{\nu}$ and \dot{P} , again pointing towards $\dot{\nu}$ as a reliable correlate of timing noise strength.

More recently, characterising the timing noise based on measured $\dot{\nu}$ values has fallen out of favor. Shannon and Cordes (2010) pointed out that although indeed the cubic term is the dominant contribution to the timing noise in an ensemble-averaged sense, in any particular realisation of the noise process significant power may lie in higher-order terms — metrics such as Δ_8 and $\sigma_z(\tau)$ will thus tend to underestimate the total contribution of timing noise to the residuals. For this reason, more recent studies have tended to return to the approach of Cordes and Helfand (1980) of using the rms timing residuals after a low-order fit as the figure of merit.

Modern Bayesian techniques based on estimation of the timing noise PSD (discussed in the previous section) modify this slightly, as the “low-order fit” is now simultaneously performed with the timing noise parameter estimation and hence is no longer straightforwardly interpreted as a weighted least-squares procedure. The conclusions of studies which adopt this approach have largely mirrored those of earlier studies — the primary factor which controls the amount of timing noise present in a pulsar is its spin-down rate (Parthasarathy et al., 2019; Lower et al., 2020). When a PSD is estimated, rather than use the rms timing residuals to assess the degree of timing noise the closely related

Table 1.1: Spearman correlation coefficients ρ and corresponding p -values calculated between the listed parameters and the rms timing residuals δt_{rms} for the 597 measurements reported by Keith et al. (2024). Scatter plots of the parameters against δt_{rms} are shown in Figure 1.3.

| Parameter | $\rho(p), A_{\text{red}} > 10^{-11}$ | $\rho(p), A_{\text{red}} < 10^{-11}$ |
|----------------------|--------------------------------------|--------------------------------------|
| σ_{RN} | 0.91 ($< 2 \times 10^{-4}$) | 0.032 (0.61) |
| A_{red} | 0.56 ($< 2 \times 10^{-4}$) | -0.034 (0.58) |
| β | 0.41 ($< 2 \times 10^{-4}$) | 0.022 (0.72) |
| $ \dot{\nu} $ | 0.32 ($< 2 \times 10^{-4}$) | -0.30 ($< 2 \times 10^{-4}$) |
| $ \ddot{\nu} $ | 0.39 ($< 2 \times 10^{-4}$) | -0.076 (0.21) |

figure of merit

$$\sigma_{\text{RN}}^2 = \int_{f_{\text{low}}}^{\infty} S(f) df = (1 \text{ yr})^{-3+\beta} \frac{A_{\text{red}}^2 f_{\text{low}}^{-(\beta-1)}}{12\pi^2 \beta - 1}, \quad (1.15)$$

is used instead, where f_{low} is the lowest frequency included in the timing noise model; typically this is the reciprocal of the observing timespan. If the timing residuals are dominated by the timing noise, then we expect the rms timing residuals δt_{rms} to be equal to σ_{RN} on average.

Figure 1.3 shows the relationship between δt_{rms} and σ_{RN} , A_{red} , β , $|\dot{\nu}|$, and $|\ddot{\nu}|$, calculated from the parameters reported by Keith et al. (2024) for 597 pulsars observed as part of the Thousand Pulsar Array project on the MeerKAT telescope (Johnston et al., 2020). Table 1.1 lists the Spearman correlation coefficients ρ and p -values between δt_{rms} and each of the parameters for those pulsars with $A_{\text{red}} > 10^{-11}$ and $A_{\text{red}} < 10^{-11}$. As expected, we have $\delta t_{\text{rms}} \approx \sigma_{\text{RN}}$ as long as the magnitude of the timing noise is not too small – the correlation is largely lost for pulsars with $A_{\text{red}} < 10^{-11}$ (indicated by the orange points), as the timing residuals become dominated by the ToA measurement noise. A similar correlation, though not as tight, is observed between δt_{rms} and A_{red} – this, again, is unsurprising, given that $\sigma_{\text{RN}} \propto A_{\text{red}}$. The spectral index β also appears to be correlated with δt_{rms} , although this is not as obvious visually – for those pulsars with $A_{\text{red}} > 10^{-11}$ we find a Spearman correlation coefficient of $\rho = 0.41$, with $p < 2 \times 10^{-4}$. For pulsars with $A_{\text{red}} < 10^{-11}$ we find $\rho = 0.02$, $p = 0.72$ – i.e. there is no significant correlation. It is not obvious where this correlation comes from – β generally describes the “kind” of timing noise in the pulsar, e.g. random walks in phase, frequency, and frequency derivative produce residual PSDs with $\beta = 2, 4$, and 6 respectively. Previous works based on both random walks and PSDs have typically not noted any correlation between β and the timing noise strength (Cordes and Downs, 1985; D’Alessandro et al., 1995; Parthasarathy et al., 2019; Lower et al., 2020).

The bottom two panels of Figure 1.3 show the scaling of δt_{rms} not with properties of the measured timing noise PSD but with recovered spin parameters of the pulsar – $|\dot{\nu}|$ and $|\ddot{\nu}|$. In both cases the $A_{\text{red}} > 10^{-11}$ sample shows significant positive correlation (see Table 1.1), as expected from the studies discussed earlier in this section. Curiously, $|\dot{\nu}|$ shows a significant negative correlation with δt_{rms} in the $A_{\text{red}} < 10^{-11}$ sample. This is likely to be related to the correlation between $|\dot{\nu}|$ and the average ToA uncertainty ($\rho = -0.41$, $p < 2 \times 10^{-4}$), which dominates the value of δt_{rms} for low A_{red} .

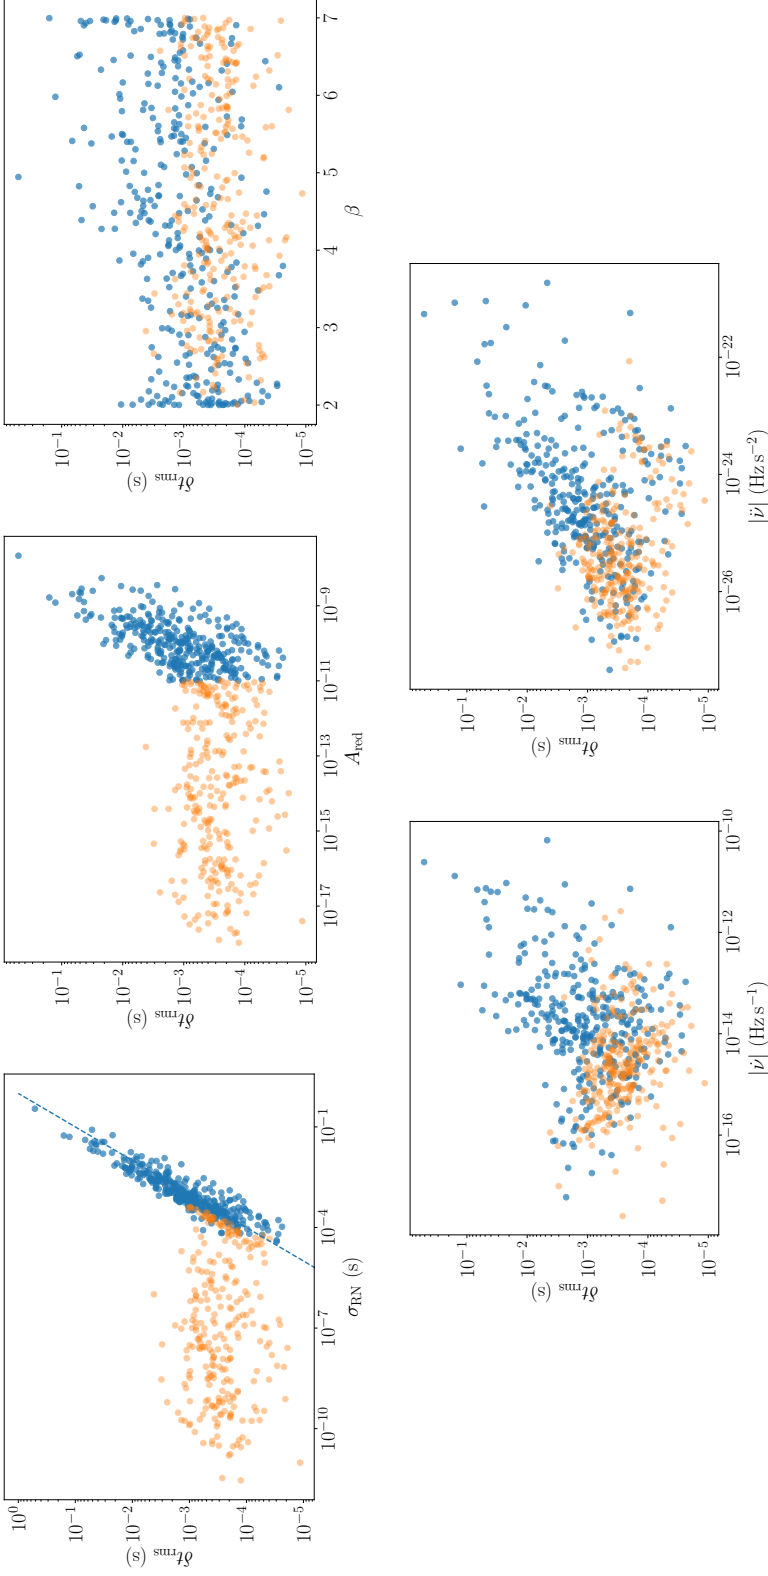


Figure 1.3: Scatter plots showing the relationship between the rms timing residuals δt_{rms} and inferred quantities related to the timing noise model (*top row*) and the spin evolution of the pulsar (*bottom row*), for 597 pulsars timed as part of the MeerKAT Thousand Pulsar Array project (Keith et al., 2024), as discussed in Section 1.4.2. The dashed line in the top left panel indicates $\delta t_{\text{rms}} = \sigma_{\text{RN}}$. The data are divided according to the value of A_{red} , with $A_{\text{red}} > 10^{-11}$ indicated by the blue points and $A_{\text{red}} < 10^{-11}$ by the orange. Spearman correlation coefficients between δt_{rms} and the other parameters are given in Table 1.1.

1.4.3 | Physical theories of timing noise

The underlying mechanism which produces timing noise is not known, and is largely outside the scope of this thesis. Here we briefly touch on some of the theories put forward in the literature which give a physical account of timing noise.

One of the earliest theories involved “starquakes” (Pines and Shaham, 1972) — small rearrangements of the solid crust, which is under significant strain, could produce frequency jumps of either sign, which produces a kind of “wandering” in the rotation of the star. However, frequency jumps alone cannot account for the full range of timing noise phenomenology which is observed (Cordes and Helfand, 1980; Cordes et al., 1988), and this model is correspondingly disfavoured as a comprehensive explanation of timing noise. We discuss starquakes further in Section 1.5.3 as a possible component of the mechanism behind pulsar glitches.

While the surface does not suffice, the stellar interior provides a promising avenue towards understanding the nature of timing noise. Anderson and Itoh (1975) pointed out that the neutron superfluid component could explain the apparent “restlessness” of pulsar spins through the stochastic motion of vortices, which are pinned to sites in the crustal lattice but are at the same time driven outward by the spin down of the star, and the stochastic unpinning and repinning of vortices imparts a time-varying torque on the crust, which is observed as timing noise. This “vortex creep” model has been refined and expanded by later authors to make predictions regarding the power spectrum of the observed noise fluctuations (Alpar et al., 1986; Jones, 1990a). Alternatively, turbulence in the superfluid component has been invoked as a possible explanation (Greenstein, 1970; Peralta, 2006; Link, 2012; Melatos and Link, 2014; Khomenko et al., 2019). Melatos and Link (2014) presented a model of timing noise due to superfluid turbulence and derived a prediction for the timing noise power spectrum, which includes a spectral turnover below a certain cutoff frequency. Goncharov et al. (2020) searched for the presence of such a turnover using both a generic broken power-law model and the specific model of Melatos and Link (2014) in seven millisecond pulsars and found that none of the pulsars significantly preferred a turnover model, although their analysis is sensitive to a turnover only within a certain frequency range dictated by the cadence of observations and the total length of the dataset.

The final category of possible timing noise models involve processes external to the stellar bulk. One particularly interesting class of models in this regard involve variable magnetospheric torques (Cheng, 1987b; Cheng, 1987a; Lyne et al., 2010). Lyne et al. (2010) reported correlations in six pulsars between the spin-down rate and variations in the pulse shape, establishing a connection (if not a direct causal link) between the magnetospheric state of the pulsar and timing noise. Other possibilities include perturbations due to planets (Kerr et al., 2015; Parthasarathy et al., 2019) or asteroid belts (Cordes and Shannon, 2008), pulse profile variation (Brook et al., 2016), and, for accreting pulsars, variability in the accretion torque (Baykal and Oegelman, 1993; Bildsten et al., 1997; Mukherjee et al., 2018; Melatos et al., 2023; O’Leary et al., 2024b).

1.5 | Glitches

For the most part, the spin frequencies of pulsars are on a more-or-less steady downward trend, as they lose their energy and angular momentum to radiation. Timing noise perturbs this trend slightly, but as discussed in Section 1.4, the magnitude of this effect is typically small and unbiased. There is, however, a more dramatic class of timing irregularities observed in a subset of pulsars – a sudden, impulsive increase in the spin frequency known as a *glitch*. Glitches are relatively rare – at the time of writing, the Jodrell Bank Observatory glitch catalogue⁸ (Basu et al., 2022) lists 671 glitches in 224 pulsars – that is, approximately 7% of all known pulsars have had at least one glitch reported. The vast majority of glitches occur in canonical pulsars – there are only two known glitches in millisecond pulsars (Cognard and Backer, 2004; McKee et al., 2016).

1.5.1 | Observational properties

The effect of a glitch on the rotational phase of the pulsar is relatively simple, typically being parametrised as follows:

$$\phi_g(t) = \Delta\phi + \Delta\nu(t - t_g) + \frac{1}{2}\Delta\dot{\nu}(t - t_g)^2 - \sum_i \tau_i \Delta\nu_i \exp[-(t - t_g)/\tau_i]. \quad (1.16)$$

Contained in this parametrisation we have a permanent phase jump $\Delta\phi$, a permanent frequency jump $\Delta\nu$, a permanent change in the frequency derivative $\Delta\dot{\nu}$, and some number (often zero, rarely more than two) of exponentially decaying components, each specified by a frequency increment $\Delta\nu_i$ and a decay timescale τ_i . The phase jump $\Delta\phi$ is not typically regarded as physical, being instead a stand-in for errors in the value of t_g , which is frequently not well-constrained⁹.

Although the limited number of events makes population-level inferences challenging, there are some overall trends which are by now well-established. There is strong evidence for two distinct populations in the glitch magnitude distribution, consisting of a broad component with $\Delta\nu$ distributed between 1×10^{-9} Hz and 1×10^{-5} Hz, and a narrow component centred around $\Delta\nu \sim 1.8 \times 10^{-5}$ Hz (Espinoza et al., 2011; Ashton et al., 2017; Fuentes et al., 2017; Basu et al., 2022), as shown in Figure 1.4. The latter component is associated with “Vela-like” pulsars (pulsars having similar inferred characteristic ages and magnetic field strengths to the Vela pulsar). Glitch activity appears to be significantly correlated with the spin-down behaviour of the pulsar, with higher spin-down rates or smaller characteristic ages associated with greater cumulative glitch activity [though note that this correlation may break down for very young pulsars, with characteristic ages $\lesssim 10^3$ yr, which has been tentatively linked to higher temperatures preventing the accumulation of sufficient stress to produce large glitches (McKenna and Lyne, 1990; Basu et al., 2022)].

For those individual pulsars which are sufficiently well-monitored and prone to glitching, the time series of glitch events is also a potentially rich source of insight. Although

⁸jb.man.ac.uk/pulsar/glitches/gTable.html

⁹Although there is the possibility of changes in the emission characteristics of the pulsar at the time of the glitch, which could manifest as a step change in the phase.




Figure 1.4: Histogram of glitch sizes recorded in the JBO glitch catalogue (Basu et al., 2022).

the majority of pulsars with sufficiently large glitch samples exhibit exponentially distributed waiting times and power-law distributed sizes, certain pulsars glitch quasiperiodically, and have strongly peaked size distributions (Melatos et al., 2008; Howitt et al., 2018; Fuentes et al., 2019). Independent of the detailed microphysics underlying glitches, it is tempting to think of the star as building up some kind of “stress” which is released in the form of a glitch. The details of this stress-relax process — how the stress accumulates, what causes it to be released, how much is released in a given event — have implications for the statistics of the glitch time series in a given pulsar. For example, if the star only glitches when a fixed stress threshold is exceeded, then we would expect to see a correlation between the size of a glitch and the waiting time until the next glitch (a *forward* correlation). In fact correlations of this type are rare (Wang et al., 2000; Melatos et al., 2018), with the only clear example being the isolated X-ray pulsar PSR J0537–6910, which glitches at a rate of $\sim 3.2 \text{ yr}^{-1}$ (Ferdman et al., 2018) and displays a strong correlation between the size of a glitch and the forward waiting time (Middleitch et al., 2006; Antonopoulou et al., 2018; Ferdman et al., 2018). However, this lack of observed correlation does not falsify the basic picture of glitching behaviour as a stress-relax process. Fulgenzi et al. (2017) introduced a model of glitching behaviour as a state-dependent Poisson process — i.e., a Poisson process where the rate is dependent on the state of the system (the amount of accumulated stress). This meta-model can reproduce the observed behaviour of many glitching pulsars and makes falsifiable predictions about size–waiting time correlations which may be observed with larger samples (Carlin and Melatos, 2019a; Carlin and Melatos, 2019b). This meta-modelling approach has also been extended to include stochastic stress accumulation (Carlin and Melatos, 2020) and history dependence (Carlin and Melatos, 2021).

Even non-detections of glitches can be constraining. For example, Millhouse et al. (2022) studied the scaling of the glitch rate in the pulsar population as a function of




Figure 1.5: Timing residuals for the pulsar PSR J0908–4913, using observations from the UTMOST pulsar timing programme. A glitch is present in the data with size $\Delta f \approx 2 \times 10^{-7}$ Hz, but is not included in the timing model. A clear cuspy feature is visible at the location of the glitch.

certain pulsar parameters. Included in their analysis were a set of pulsars which had not been observed to glitch but for which glitches above a certain size had been systematically ruled out (see Chapter 3 of this thesis). They found that including these non-glitching pulsars significantly shifted the inferred population parameters — non-detections provide valuable information, as long as the detection procedure is well-understood.

1.5.2 | Glitch detection

Glitches are most often detected from by-eye inspection of the timing residuals — if a timing model without a glitch is fit to data containing a glitch, one tends to see a characteristic cuspy feature at the glitch epoch, as shown in Figure 1.5. By-eye detection has the obvious advantage that it is trivial to implement, and for large glitches it is entirely sufficient. However, the distribution of glitch sizes extends down to small amplitudes where the changes in spin parameters are comparable to the perturbations generated by timing noise. The low-amplitude end of the observed glitch distribution varies significantly among pulsars. Figure 1.6 shows the distribution of minimum measured fractional glitch size in every pulsar with at least two glitches recorded, along with the minimum, median and maximum 90% confidence upper limits on the size of undetected glitches derived from the analyses described in Chapter 5. For the purposes of having an accurate statistical picture of the glitch population, two questions in particular are of interest: what sorts of glitches might have been missed, and is every detected glitch genuine?

Can by-eye inspection of timing residuals reliably determine which features in the residuals are due to glitches, and which are simply part of the long-term timing noise




Figure 1.6: Distribution of minimum observed fractional glitch amplitude among pulsars with at least two glitches recorded. The dotted lines indicate the minimum (2.4×10^{-12}) and maximum (1.3×10^{-7}) 90% confidence upper limits on the fractional size of undetected glitches derived from the analyses described in Chapter 5, while the solid line shows the median 90% upper limit (2.1×10^{-9}).

behaviour? Clearly there is a point at which a glitch is undetectable in a given dataset, but where is the crossover? Janssen and Stappers (2006) reported glitches with sizes as small as $\Delta\nu/\nu = 10^{-11}$, and performed synthetic data tests to demonstrate that they could reliably recover injected glitches of that magnitude when inspecting the synthetic residuals based on the dataset in which the smallest glitches were found. However, verifying successful recovery of injected glitches leaves open the question of false positives – timing noise events which are mis-identified as small glitches. It is also not scalable – the limits of by-eye detection depend on the intrinsic characteristics of the pulsar as well as the characteristics of a given dataset, and to perform injection studies of this sort might be plausible for a handful of datasets, but covering hundreds of datasets each covering many years of timing baseline has a significant cost in person-hours.

Automated methods of glitch detection provide a possible way forward – if input from a human analyst is not required, performing a large number of injection studies to assess false positive and negative rates for different pulsar and glitch parameters becomes much more feasible (provided the automated method is not too computationally expensive). The results are also uniform in a way which is hard to guarantee for by-eye detections: humans are subject to tiredness, distraction, etc., which may affect their performance as a glitch detector in ways which are difficult to control for. Of course, no glitch detector is perfect – false positives and negatives can not be avoided, and as such statements about glitch detection and catalogue completeness are necessarily statistical in nature.

A number of automated glitch detection methods have been proposed and implemented. For the purposes of this thesis the most notable is the hidden Markov Model-based method introduced by Melatos et al. (2020), which is described in Section 1.10 and

applied in Chapters 2, 3, 4 and 5 of this thesis. Espinoza et al. (2014) described an automated method based on goodness-of-fit measurements to moving windows of ToAs. They used this method to argue that the absence of observed small glitches in the Crab pulsar was a real effect rather than an artefact of detectability limits, and later analyzed timing data from the Vela pulsar to find two previously unknown small glitches (Espinoza et al., 2021). However, they did not evaluate the performance of the detector explicitly, and the distinction between a genuine glitch and a “glitch candidate” returned by the detector is not precise. Yu and Liu (2017) discussed the detection probability of glitches using synthetic datasets with known injected glitch parameters and analyzing the accuracy of the glitch parameters recovered using `TEMPONEST`. They used these results to argue that the glitch catalogue reported by Yu et al. (2013) is likely to be complete or almost-complete. However, they did not discuss how to evaluate detection in the absence of *a priori* known glitch parameters – their detection criteria involve knowledge of the injected glitch parameters. Singha et al. (2021a) described a glitch detection pipeline which operates with low latency at the Ooty Radio Telescope, and is based on detection of outlier ToAs combined with a low-order polynomial regression to account for the presence of timing noise. They demonstrate the effectiveness of this pipeline using synthetic data tests for a range of timing noise parameters, and a glitch in PSR J0742–2822 has been reported which was detected by this system (Grover et al., 2022). At this stage, however, this pipeline has not been used to derive upper limits on the size of missed glitches and hence assess the completeness of the glitch catalogues derived from the application of this method.

1.5.3 | Physical theories of glitches

As with timing noise, the physical processes behind glitches remain unknown, but in the 55 years since their discovery (Radhakrishnan and Manchester, 1969; Reichley and Downs, 1969) they have received significant theoretical attention – see Haskell and Melatos (2015) for a comprehensive review of proposed glitch mechanisms. Here we give only a brief overview of the landscape of models.

Given that most glitches in rotation-powered pulsars do not seem to be associated with radiative changes (though see Weltevrede et al., 2011; Palfreyman et al., 2018 for exceptions), the focus has largely centred around mechanisms involving the crust or interior of the star, similarly to the timing noise models discussed in the Section 1.4.3.

Mechanical failures of the crystalline crust – “star quakes” – are an attractive option, and were proposed as a possible explanation soon after the first glitches were detected (Ruderman, 1969). The essential idea is that as the neutron star spins down, the equilibrium shape deviates slightly from the shape of the rigid crust, and this causes a build-up of stress which is released in a sudden reorganisation of the crystal lattice, leading to a small reduction in the moment of inertia and hence a spin-up event. However, pulsars which exhibit large, frequent glitches (e.g. the Vela pulsar) cannot be accommodated in this model – there is no hope of building up sufficient stress quickly enough to explain the observed glitching activity in such pulsars (Baym and Pines, 1971; Rencoret et al., 2021). Though glitches cannot be explained through starquakes alone, crustal failures may nevertheless play an important role as a glitch triggering mechanism, driving internal processes which ultimately lead to transfer of angular momentum to the crust

(Ruderman, 1976; Alpar et al., 1994; Link and Epstein, 1996; Middleditch et al., 2006; Bransgrove et al., 2020). Glitches thus provide a potential window into the microphysical properties of the crust (Horowitz and Kadau, 2009; Chugunov and Horowitz, 2010), as well as its global behaviour under strain (Franco et al., 2000; Giliberti et al., 2019; Giliberti et al., 2020; Kerin and Melatos, 2022). Mechanical failure of the crust is also of interest from the perspective of continuous gravitational wave emission, as the formation of “mountains” – small non-axisymmetries in the mass distribution – are a candidate emission mechanism (Ushomirsky et al., 2000; Johnson-McDaniel and Owen, 2013; Kerin and Melatos, 2022; Gangwar and Jones, 2024), see Section 1.7.1.

The standard picture, then, is one of angular momentum transfer from the superfluid component of the star to the crust (Anderson and Itoh, 1975). The angular momentum of a rotating superfluid is stored in discrete vortices, each of which carries a fixed quantum of angular momentum. For the superfluid to decelerate, these vortices must migrate outwards and ultimately leave the system. However, the vortices may pin to the crust or to magnetic flux tubes in the interior (see Section 1.1.1), preventing them from moving outward and hence preventing the superfluid component from decelerating. The crust, however, is decelerated due to the electromagnetic braking torque, and so differential rotation develops between the crust and the superfluid. A glitch occurs when, for one reason or another, some of the pinned vortices unpin and are allowed to migrate outwards, in the process transferring angular momentum to the crust. The triggering event which causes the large-scale unpinning is unclear, and models abound, including starquakes (see above), hydrodynamic instabilities (Andersson et al., 2003; Mastrano and Melatos, 2005; Glampedakis and Andersson, 2009), or collective motion due to knock-on effects of local unpinning events (“avalanches”) (Melatos et al., 2008; Warszawski et al., 2012; Warszawski and Melatos, 2013; Khomenko and Haskell, 2018; Howitt et al., 2020).

A full theory of glitches must explain more than the triggering mechanism and increase in spin frequency – it must also explain the observed relaxations. A change in the period derivative was noted in the very first glitch to be observed (Radhakrishnan and Manchester, 1969; Reichley and Downs, 1969), and Baym et al. (1969b) interpreted this as a quasi-exponential relaxation back to the pre-glitch spin-down state, with a timescale of order years, and hence argued that this was a signature of the presence of a superfluid component coupled to the crust. Although the exponential nature of the decay was not clear in 1969, a subsequent measurement of a glitch in the Crab pulsar by Boynton et al. (1972) provided an example of a much faster exponential recovery, with a timescale on the order of days. Alpar et al. (1984b) developed a theory of glitch relaxation in terms of the thermally driven “creep” of vortices outward, which recouples the crust and superfluid component on a timescale which is dependent on the internal temperature of the star but is generally within the observed range of relaxation timescales for plausible physical parameters. Alpar et al. (1984a) applied this theory successfully to four large glitches in the Vela pulsar, and obtained an estimate of the internal temperature of Vela of 10^7 K. Vortex creep theory has been significantly extended, refined, and applied since this initial work (Alpar et al., 1989; Link et al., 1993; Alpar et al., 1996; Link, 2006; Link, 2014; Akbal et al., 2017; Gügercinoğlu et al., 2022).

Other post-glitch relaxation mechanisms have been proposed. Motivated by the extremely large glitch observed in PSR B0355+54 (Lyne, 1987), Jones (1990b) proposed a model where a region of the superfluid component in the crust is only weakly pinned

and thus the vortices are in near-corotation with the superfluid, and showed that the relaxation process could be explained via the interaction between this weakly pinned component and the crystal lattice. Haskell et al. (2012) elaborated on this essential picture with a hydrodynamical model which successfully reproduces the post-glitch behaviour observed in the Vela pulsar for a physically plausible set of parameters. The presence of turbulence in the superfluid component significantly modifies the post-glitch behaviour, as the coupling of the superfluid component to the lattice is dependent on the vortex structure, which becomes tangled in the turbulent case (Peralta et al., 2006; Drummond and Melatos, 2017; Mongiovì et al., 2017). Haskell et al. (2020) have argued that the observed post-glitch behaviours in PSR J0537–6910 and the Vela pulsar point towards the presence of a turbulent superfluid component.

1.6 | Gravitational waves in general

The theory of general relativity (GR) describes spacetime as a four-dimensional manifold whose curvature depends on the distribution of matter and energy via the Einstein Field Equations (in units where $G = c = 1$):

$$G_{\mu\nu} = 8\pi T_{\mu\nu}, \quad (1.17)$$

where $G_{\mu\nu}$ is the Einstein tensor which describes the spacetime curvature and $T_{\mu\nu}$ is the stress-energy tensor which describes the distribution of matter and energy¹⁰. One of the most surprising features of GR is that it supports wave-like vacuum ($T_{\mu\nu} = 0$) solutions, known as gravitational waves (GWs). The existence of such solutions has been known since the early days of GR (Einstein, 1918), but for several decades it was not clear that they transported energy or otherwise had any physical effect – one must be careful in GR to distinguish physical results from artifacts of one particular coordinate system or another, and many (including Einstein himself) worried that the apparent wave-like solutions might simply be an illusion. However, in 1957 Feynman put forth a simple thought experiment known as the “sticky bead argument” which established that indeed gravitational waves could have physical effects, and so one could in principle build a detector to measure them (Feynman, 1957)¹¹.

Observational evidence for the existence of GWs came nearly two decades later, when Hulse and Taylor (1975) discovered the pulsar PSR J1915+1606, which is in a binary system with another neutron star. The high precision afforded by pulsar timing allowed for measurement of the orbital decay of this system due to gravitational radiation, and the rate of decay is in good agreement with the GR prediction (Taylor et al., 1979; Weisberg et al., 2010). Efforts at direct detection were also occurring around this time, most notably by Weber, who pioneered the use of large metal bars with a resonant frequency on the order of 1 kHz, which would be “rung up” by a passing gravitational wave whose frequency matched the resonant frequency of the bar. Although Weber claimed evidence for the detection of a GW signal coming from the direction of the Galactic centre

¹⁰Strictly speaking there is a further term on the left-hand side, $\Lambda g_{\mu\nu}$ where Λ is the “cosmological constant” and $g_{\mu\nu}$ is the spacetime metric. In this work we will assume $\Lambda = 0$ always.

¹¹In that same year, Bondi (1957) gave a similar argument, and Weber and Wheeler (1957) also argued strongly for the physicality of gravitational waves.

(Weber, 1969; Weber, 1970), the discovery could not be independently verified (Thorne, 1980), and is not regarded today as representing a true detection.

Successful direct detection of gravitational waves came on September 14, 2015 at 09:50:45 UTC, when the Laser Interferometer Gravitational-Wave Observatory (LIGO) project detected a signal, dubbed GW150914, from the merger of two $\sim 30 M_\odot$ black holes at a distance of approximately 400 Mpc (Abbott et al., 2016b). In the years since, LIGO (and later the Virgo and KAGRA detectors) have amassed a catalogue of approximately 100 signals from mergers of compact binaries (Abbott et al., 2019d; Abbott et al., 2021a; Abbott et al., 2023a), including GW170817, a nearby merger of two neutron stars which was also observed across the entire electromagnetic spectrum (Abbott et al., 2017c; Abbott et al., 2017b).

1.6.1 | A short mathematical discussion

In this section we briefly discuss the mathematical formalism of gravitational waves, following Maggiore (2007).

To begin with, we linearise about the flat-space metric $\eta_{\mu\nu}$, i.e.

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \quad |h_{\mu\nu}| \ll 1. \quad (1.18)$$

It is useful to define the trace-reversed metric perturbation,

$$\bar{h}_{\mu\nu} = h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}h, \quad (1.19)$$

where $h = \eta^{\mu\nu}h_{\mu\nu}$. It is also useful to impose a gauge condition (the Lorentz gauge),

$$\partial^\nu \bar{h}_{\mu\nu} = 0. \quad (1.20)$$

With this in hand, one can show that the equations of motion reduce to

$$\square \bar{h}_{\mu\nu} = -16\pi T_{\mu\nu}, \quad (1.21)$$

where $\square = \partial^\mu \partial_\mu$. This is a wave equation for $\bar{h}_{\mu\nu}$! The waves are sourced by $T_{\mu\nu}$, the stress-energy distribution.

Far from the source (i.e. in the region where $T_{\mu\nu} = 0$), we can greatly simplify the form of the metric by adopting the *transverse-traceless* (TT) gauge, which further imposes

$$\bar{h} = 0, \quad (1.22)$$

$$h^{0\mu} = 0, \quad (1.23)$$

$$h_i^i = 0, \quad (1.24)$$

$$\partial^j h_{ij} = 0, \quad (1.25)$$

where \bar{h} is the trace of $\bar{h}_{\mu\nu}$ and this first condition implies that $\bar{h}_{\mu\nu} = h_{\mu\nu}$ (hence the dropping of the bar on the subsequent gauge conditions).

Equation 1.21 admits plane wave solutions. Without loss of generality we may take the direction of propagation to be along the z axis, and in that case the plane wave solution takes the form

$$h_{ij}^{\text{TT}}(t, z) = \begin{pmatrix} h_+ & h_\times & 0 \\ h_\times & -h_+ & 0 \\ 0 & 0 & 0 \end{pmatrix}_{ij} \cos[\omega(t - z)]. \quad (1.26)$$




Figure 1.7: Representative noise amplitude spectral density curves for the Hanford (H1) LIGO detector during the first observing run (O1) (Sigg, 2016) and second half of the third observing run (O3b) (Goetz, 2021).

The h_+ and h_\times are the amplitudes of the “plus” and “cross” polarisations, so called because of their effect on a ring of freely-falling test masses in the plane transverse to propagation.

1.6.2 | Ground-based gravitational wave detectors

The GW-focused parts of this thesis are concerned with the analysis of data taken by ground-based laser interferometric gravitational wave detectors, particularly the LIGO detectors. The LIGO collaboration operates two detectors, located in Hanford, Washington (H1) and Livingston, Louisiana (L1) in the United States of America. Both are dual-recycled Fabry-Pérot Michelson interferometers with two 4 km arms (Buikema et al., 2020). The passage of a gravitational wave through the detector causes a differential change in the length of the two arms, which is ultimately converted into an estimate of the dimensionless strain $h(t) = \Delta L/L$ (Sun et al., 2020b).

The era of Advanced LIGO began in earnest with the first Observing Run (O1), which ran from September 2015 to January 2016 (Abbott et al., 2016a). Subsequently there have been two more complete observing runs — O2 ran from November 2016 to August 2017 (Abbott et al., 2021c), and O3 ran from April 2019 to March 2020, divided into O3a and O3b (Abbott et al., 2023b). The fourth observing run is ongoing at the time of writing. The sensitivity of the detectors has steadily improved over this period — Figure 1.7 shows a comparison of the noise amplitude spectral density (ASD) between O1 and O3b in the Hanford detector. In the most sensitive frequency range of the O3b noise curve, around 250 Hz, the ASD is smaller in O3b by a factor of approximately 1.7.

Besides LIGO, there are three other active ground-based interferometric detectors: Virgo, located near Pisa, Italy (Acernese et al., 2015); the Kamioka Gravitational Wave Detector (KAGRA) located near Toyama, Japan (Kagra Collaboration et al., 2019); and GEO600 located near Hannover, Germany (Lück and GEO600 Team, 1997). While we do not use data from these detectors in this work, the LIGO, Virgo and KAGRA collaborations together operate as a global gravitational wave detection network (Brady et al., 2022).

1.7 | Continuous gravitational waves

Broadly speaking, gravitational wave signals can be divided into four quadrants, according to whether they are short-duration or long-duration, and modelled or unmodelled. To date, all direct detections of gravitational waves fall into the short-duration, modelled quadrant – these are compact binary coalescence signals, produced by final stage of inspiral and subsequent merger of neutron stars and black holes (Abbott et al., 2016b; Abbott et al., 2017b; Abbott et al., 2023a). These signals are typically in the frequency band accessible to ground-based detectors (roughly 20 Hz–4 kHz) for a minute or less, and the form of the signal can be predicted with high accuracy using a combination of numerical relativity and post-Newtonian approximations (Blanchet, 2014).

In this thesis, we will be concerned with the (long-duration, modelled) quadrant – commonly referred to as “continuous wave” (CW) signals. These signals are typically assumed to last for timescales much longer than the timespan of any given dataset, i.e. approximately one year¹², and the form of the signal is, to a first approximation, exceptionally simple: it is essentially a sinusoid with near-constant frequency.

1.7.1 | Producing continuous gravitational waves

There are multiple physical mechanisms which may produce detectable CW signals, from the (relatively) mundane to the exotic. In the LIGO frequency band, the most mundane (and so in some sense the most promising) mechanisms rely on neutron stars to provide the necessary contribution to $T_{\mu\nu}$.

The simplest mechanism by which a neutron star might generate continuous gravitational radiation is via a “mountain” – a non-axisymmetric mass distribution on the neutron star surface. Such a mountain may be formed in a variety of ways, depending on the history of the star. A neutron star which is accreting material from a companion will funnel this matter towards the magnetic poles, creating a mountain of accreted material (Melatos and Payne, 2005; Priymak et al., 2011; Barboza Rossetto et al., 2024) or producing temperature gradients which deform the crust (Bildsten, 1998; Ushomirsky et al., 2000; Hutchins and Jones, 2023). In the isolated case, mountains may arise due to deformations of the crust as the neutron star cools and spins down (Fattoyev et al., 2018; Kerin and Melatos, 2022; Gangwar and Jones, 2024), or from strains due to the magnetic field (Palomba, 2001; Ioka, 2001; Lasky and Melatos, 2013). The mountain size is typically

¹²Signals which are transient on timescales \sim hours–months are also considered in some analyses and are generally categorised as CW signals (e.g. Abbott et al., 2019b; Abbott et al., 2022a), but they are not the focus of this thesis.

characterised by the ellipticity ϵ , given by

$$\epsilon = \frac{|I_{xx} - I_{yy}|}{I_{zz}}, \quad (1.27)$$

where I_{zz} is the moment of inertia about the axis of rotation (assumed to be one of the principal axes), and I_{xx} and I_{yy} are the moments of inertia about the other two principal axes. The strain amplitude is then given by (Jaranowski et al., 1998)

$$h_0 = \frac{4\pi^2 G \epsilon I_{zz} f_{\text{GW}}^2}{c^4 d} \quad (1.28)$$

$$= 1.1 \times 10^{-26} \left(\frac{\epsilon}{10^{-6}} \right) \left(\frac{I_{zz}}{10^{38} \text{ kg m}^2} \right) \left(\frac{d}{1 \text{ kpc}} \right)^{-1} \left(\frac{f_{\text{GW}}}{100 \text{ Hz}} \right)^2 \quad (1.29)$$

where $h_0 \sim |h_{+/\times}|$ and $f_{\text{GW}} = 2f_{\text{rot}}$ is the frequency of the emitted gravitational radiation and d is the distance to the source.

Another commonly considered mechanism for CW generation is emission via r -modes. These are fluid oscillation modes which may be excited in the interior of the neutron star, and produce CW radiation via a mass current quadrupole (Arras et al., 2003; Santiago-Prieto et al., 2012). These r -modes are of particular interest because they are *unstable* to the emission of gravitational waves (Andersson, 1998; Owen et al., 1998; Caride et al., 2019). The frequency of the CW signal is equal to approximately 4/3 the rotation frequency of the star, although the exact proportionality depends on the properties of the neutron star interior. For a fiducial neutron star model, the strain amplitude is given by (Owen, 2010)

$$h_0 = 3.6 \times 10^{-26} \left(\frac{\alpha}{10^{-3}} \right) \left(\frac{d}{1 \text{ kpc}} \right)^{-1} \left(\frac{f_{\text{GW}}}{100 \text{ Hz}} \right)^3, \quad (1.30)$$

where α is the dimensionless r -mode amplitude.

Apart from neutron stars, the other source of CW emission in the LIGO band which is commonly considered is from ultralight boson clouds around black holes. Such ultralight bosons are not known to exist, but have been proposed as a component of the dark matter in the Universe (Preskill et al., 1983). If they do exist, they can be produced in large numbers around black holes via superradiance (Zel'Dovich, 1971; Starobinskiĭ, 1973; Arvanitaki et al., 2010), after which they subsequently produce continuous gravitational radiation. The radiation is produced either by direct annihilation of the bosons into gravitons or via transitions between energy levels around the black hole, in close analogy with electromagnetic radiation due to electronic transitions observed in conventional atoms (Brito et al., 2015; Sun et al., 2020a; Abbott et al., 2022e).

1.7.2 | Detecting continuous gravitational waves

The basic signal model for CW detection is very similar to that used in pulsar timing, as discussed in Section 1.2.2. In the source frame, the phase of the signal is again described by a Taylor series about a reference time t_0 (see equation 1.4), and the transformation into the detector frame introduces corrections which depend on various extrinsic factors, such as sky position and possible binary motion. However, unlike the pulsar timing case where we have highly precise but sparse measurements of the phase in the form of times

of arrival, in the CW case we are faced with a continuous stream of data containing a weak sinusoidal signal buried in the noise. Despite the variety of mechanisms for generating CW signals discussed in the previous section, the ultimate form of the signal is not essentially different (except that in some cases signal power may be present at more than one harmonic of the rotation frequency), and so most detection strategies are quite agnostic with respect to the underlying mechanism. Here we will focus on a matched-filtering approach introduced by Jaranowski et al. (1998), which is the basis for the search carried out in Chapter 6.

We assume the detector data $x(t)$ is a linear combination of Gaussian noise $n(t)$ and the signal

$$h(t) = A^\mu h_\mu(t), \quad (1.31)$$

where $\mu = \{1, 2, 3, 4\}$ and the repeated index is summed over. The A^μ are the “amplitude parameters”, time-independent numbers depending on the geometry of the system, the strain amplitude h_0 , and the initial phase Φ_0 . The $h_\mu(t)$ are schematically of the form

$$h_\mu(t) \sim a_\mu(t) \sin \Phi(t), \quad (1.32)$$

where $a_\mu(t)$ is a slowly-varying function which describes the detector response to the incoming wave and $\Phi(t)$ is the signal phase in the detector frame. The full forms of these terms can be found in Jaranowski et al. (1998). The log likelihood function is then given by the usual matched-filtering likelihood

$$\ln \Lambda = \langle x|h \rangle - \frac{1}{2} \langle h|h \rangle = A^\mu \langle x|h_\mu \rangle - \frac{1}{2} A^\mu \langle h_\mu|h_\nu \rangle A^\nu, \quad (1.33)$$

where the inner product $\langle x|y \rangle$ is defined as

$$\langle x|y \rangle = 4\Re \int_0^\infty df \frac{x(f)y^*(f)}{S_h(f)} \approx \frac{2}{S_h(f_0)} \int_0^T dt x(t)y(t) \quad (1.34)$$

where the approximate equality holds in the case that the frequency content of $x(t)$ and $y(t)$ is contained in a sufficiently narrow band around a central frequency f_0 that the one-sided noise PSD $S_h(f)$ can be taken as constant.

The decomposition of the signal into the form in equation (1.31) turns out to be convenient because it is possible to analytically maximize the likelihood function over the A^μ , which are typically not of significant physical interest. This partially-maximized likelihood is referred to as the \mathcal{F} -statistic (Jaranowski et al., 1998). It is also possible to instead treat the problem in a Bayesian framework and *marginalise* over the amplitude parameters using a physically motivated prior – this leads instead to the so-called \mathcal{B} -statistic, although this is rarely used in preference to the \mathcal{F} -statistic because it is significantly more computationally expensive to calculate (Prix and Krishnan, 2009; Whelan et al., 2014). After eliminating the amplitude parameters by maximisation or marginalisation, there are $3 + s$ parameters which must be chosen – two to specify the sky position, f_{GW} , and the first s time derivatives of f_{GW} (similarly to the pulsar timing case, s is frequently one and rarely larger than two)¹³. If these parameters are not known *a priori* with high

¹³If binary motion is relevant then extra parameters are needed, but in this thesis we will not be concerned with this complication.

precision then they must be explicitly searched over, by constructing a template bank which contains many different choices of the unconstrained parameters, and evaluating the detection statistic for each choice.

Thus far, the signal model has been *coherent* — for a given choice of signal parameters, the signal phase $\Phi(t)$ is continuous and deterministic for the duration of the data set. While this approach is highly sensitive, it is also typically computationally prohibitive unless the signal parameters are extremely well constrained — the required number of distinct choices of signal parameters (“templates”) needed to cover a region of parameter space grows quickly with the observing timespan, and for the typical timespans of months–years a fully coherent approach is restricted to targets where the parameter space is very small, e.g. searches for CW emission from known pulsars with electromagnetically derived ephemerides (Abbott et al., 2022h). In order to render broader searches tractable, a *semicoherent* approach is often adopted instead. There are a wide variety of semicoherent approaches, but the essential common idea is to divide up the data to be analysed into smaller segments, analyse these segments coherently, and then combine the results of the coherent subanalyses. This approach trades sensitivity for reduced computational cost. The sensitivity of a semicoherent search scales like $(T_{\text{obs}}T_{\text{coh}})^{-1/4}$ where T_{obs} is the total observing time and T_{coh} is the length of the coherent segments, compared to $T_{\text{obs}}^{-1/2}$ for a fully coherent analysis (Krishnan et al., 2004)¹⁴. However, the scaling of the required number of templates is significantly reduced from T_{obs}^n with $n \sim 5$ to $N_{\text{seg}}^m T_{\text{coh}}^n$ where $N_{\text{seg}} = T_{\text{obs}}/T_{\text{coh}}$ is the number of coherent segments and $m \sim 2$ (Prix and Shaltev, 2012; Wette et al., 2018).

Semicoherent searches can also offer an additional benefit over fully coherent searches: flexibility in the face of deviations from a canonical signal model. By dividing the data into segments which are analysed separately and then combining the results, one gains the freedom to stitch together the results from a series of templates which are not identical in one or more parameters, thereby greatly expanding the family of signal models under consideration. Most often this is done by allowing the template frequency to change between timesteps — this is the idea between the various hidden Markov model-based searches which have been developed to efficiently explore signal models which include a degree of stochastic frequency evolution, as discussed further in Section 1.9.

1.7.3 | Searches for continuous waves

Since the first continuous wave searches on data from ground-based interferometers, approximately 300 searches have been published with a wide variety of techniques, data sets, and targets. A comprehensive survey of the history of this endeavour is beyond the scope of this thesis; for recent reviews see Piccinni (2022), Riles (2023), and Wette (2023). Here we aim only to give a brief overview of the landscape.

Continuous wave searches are typically divided into three broad categories, according to strategy: *targeted*, *directed*, and *all-sky*.

- *Targeted* searches look for emission from sources already observed as electromagnetic pulsars, and make the assumption that the continuous wave emission

¹⁴Strictly speaking this scaling applies in the limit where the probability of false alarm is chosen to be small and the number of coherent segments is large — see Wette (2012).

is phase-locked to the electromagnetic emission, up to an overall factor depending on the assumed CW emission mechanism (e.g. for emission by a mountain the CW phase is twice the EM phase) (Abbott et al., 2010; Nieder et al., 2019; Abbott et al., 2022h). Because the covered parameter space is essentially a single point for each target in these searches, they are both computationally cheap and highly sensitive, but of course the strong assumptions made in these searches may well not be satisfied in reality. Certain searches relax this assumption of phase locking somewhat, allowing for the frequency of the signal to deviate from the electromagnetically inferred value within a small range — these are referred to as *narrowband* searches (Abbott et al., 2019b; Ashok et al., 2021; Abbott et al., 2022a).

- *Directed* searches aim to detect emission from a neutron star but do not use an electromagnetically measured ephemeris. Most often this is because the target is a known (or assumed) neutron star which is not observed as a pulsar — typical targets are supernova remnants assumed to contain a young neutron star, or low-mass X-ray binaries where the compact object is likely to be a neutron star (see e.g. Wette et al., 2008; Whelan et al., 2015; Beniwal et al., 2021; Abbott et al., 2022f; Abbott et al., 2022c). As there is no direct information available on the neutron star’s spin frequency, these searches are necessarily broad-band, and frequently must cover a non-trivial volume in spin-down and/or binary orbit parameter space (Abbott et al., 2022c; Abbott et al., 2022f). Directed searches also encompass searches which do not target a particular putative neutron star, but “interesting” sky regions where there may be a population of neutron stars, for example globular clusters or the galactic center, both of which are very dense stellar environments which are expected to host significant populations of fast-spinning neutron stars (Aasi et al., 2013; Abbott et al., 2017a; Dergachev et al., 2019; Piccinni et al., 2020).
- *All-sky* searches, as the name suggests, are the most “uninformed” searches — they aim to find CW emission from unknown neutron stars which could be located anywhere on the sky. As with directed searches, these searches are typically broad-band (although not always, see Wette et al., 2021 for an example of a narrowband all-sky search), and may or may not include a search over binary orbital parameters in addition to frequency and spin-down (Abbott et al., 2008a; Abbott et al., 2022b; Abbott et al., 2021b; Dergachev and Papa, 2021; Dergachev and Papa, 2023). These searches are typically extremely computationally intensive, and relatively insensitive in terms of the limits placed on the characteristic strain. However, only a small fraction of the neutron stars in our galaxy are visible as pulsars, LMXBs, etc., (Sartore et al., 2010) and even with reduced sensitivity compared to more informed searches, it is valuable to perform all-sky searches which allow for sensitivity to the rest of the galactic population, particularly those unseen neutron stars which may be close by.

Within each of these categories a wide variety of approaches have been developed over the years, both in the tools used for analysis and in the choices made regarding the signal model and parameter space covered. Table 1.2 lists a cross-section of searches among the three categories described above, focussing on recent searches and those using hidden Markov models and the Viterbi algorithm (see Section 1.9).

Table 1.2: Details of a cross-section of published CW searches, covering the three categories discussed in Section 1.7.3. Here we focus on recent searches and searches directly relevant to this thesis, i.e. those using some flavour of Viterbi and those targeting particular sky regions. See Wette (2023) for a comprehensive discussion of past CW searches. Upper limits on h_0 are conditional on the details of the analysis and should not be compared directly between searches. If f_{GW} is listed as “Various” the search comprises a number of distinct narrow frequency ranges informed by electromagnetic observations of the targets’ frequencies. Abbreviations used — Gal. Cen: Galactic Center; LC: loosely coherent; SNR: supernova remnant; AMXP: accreting millisecond X-ray pulsar; NB: narrowband; NS: neutron star. References for methods: ^a Dergachev (2010); ^b Jaranowski et al. (1998); ^c Piccinni et al. (2020); ^d Suvorova et al. (2016); ^e Suvorova et al. (2017); ^f Pitkin et al. (2017); ^g Astone et al. (2012); ^h Astone et al. (2014); ⁱ Krishnan et al. (2004); ^k Bayley et al. (2019); ^l Covas and Sintees (2019); ^m Dergachev and Papa (2020).

| Category | Target(s) | Method | Dataset | f_{GW} (Hz) | Best h_0 upper limit | Reference | |
|-------------|---------------------------------------|---------------------------------------|--------------------------|----------------------------|--------------------------------|--------------------------------|--------------------------------|
| Directed | Gal. Cen. & Ter 5 | LC PowerFlux ^a | O1 | 475–1500 | 2.9×10^{-25} @ 550 Hz | Dergachev et al. (2019) | |
| | NGC 6544 | \mathcal{F} –statistic ^b | S6 | 92.5–675 | 6.7×10^{-25} @ 173 Hz | Abbott et al. (2017a) | |
| | Orion spur | LC PowerFlux ^a | S6 | 50–1500 | 6.3×10^{-25} @ 169 Hz | Aasi et al. (2016) | |
| | Gal. Cen. | FrequencyHough ^c | O3 | 10–2000 | 7.6×10^{-26} @ 142 Hz | Abbott et al. (2022h) | |
| | Scorpius X-1 | Viterbi ^{d,e} | O3 | 60–500 | 6.2×10^{-26} @ 256 Hz | Abbott et al. (2022f) | |
| | Young SNRs | Viterbi ^d | O3a | Various | 2.6×10^{-25} @ 172 Hz | Abbott et al. (2021d) | |
| | Targeted | AMXPs | Viterbi ^{d,e} | O3 | Various | 4.7×10^{-26} @ 218 Hz | Abbott et al. (2022g) |
| | | Known pulsars | Various ^{b,f,g} | O2+O3 | Various | 4.7×10^{-27} @ 103 Hz | Abbott et al. (2022e) |
| | | Known pulsars (NB) | Various ^{b,g} | O3 | Various | 2.3×10^{-26} @ 364 Hz | Abbott et al. (2022a) |
| | | All-sky | Isolated NS | Various ^{b,h,i,k} | O3 | 10–2048 | 1.1×10^{-25} @ 112 Hz |
| Binary NS | BinarySkyHough ^l | | O3a | 50–300 | 2.4×10^{-25} @ 150 Hz | Abbott et al. (2021b) | |
| Isolated NS | Falcon ^m | | O3a | 500–1000 | 1.0×10^{-25} @ 523 Hz | Dergachev and Papa (2023) | |
| Isolated NS | \mathcal{F} –statistic ^b | | O3 | 20–800 | 8.1×10^{-26} @ 203 Hz | Steltner et al. (2023) | |

1.8 | Hidden Markov models

Hidden Markov models (HMMs) provide a framework for analysing systems where the evolution of the system is memoryless (*Markovian*), and our observations of the system yield probability distributions over the system state, rather than direct knowledge of its state (the state is *hidden*). In this section we will introduce the formalism surrounding HMMs in detail, and discuss their application to two problems which are to be discussed in the rest of this thesis: the detection and characterisation of pulsar glitches, and the detection of continuous gravitational waves.

1.8.1 | Formalism

Formally, an HMM consists of the following components:

- A set of possible hidden states $\mathcal{Q} = \{q_1, q_2, \dots, q_{N_Q}\}$. In our case \mathcal{Q} will always be finite – extensions to continuous \mathcal{Q} are possible but not used here.
- A set of discretized timesteps $\mathcal{T} = \{t_0, t_1, \dots, t_{N_T}\}$. Observations are made at t_1, \dots, t_{N_T} . The timestep t_0 is a special timestep which occurs before any measurement is made.
- A set of possible observations \mathcal{O} . There is little restriction on \mathcal{O} – it may be finite or infinite, discrete or continuous.
- A transition probability function

$$A(q_i, q_j; t_k) = \Pr[q(t_{k+1}) = q_i \mid q(t_k) = q_j]. \quad (1.35)$$

This function is not always an explicit function of t_k .

- An emission likelihood function

$$L(o, q_i; t_k) = \Pr[q(t_k) = q_i \mid o(t_k) = o]. \quad (1.36)$$

- A prior distribution on the hidden states

$$\Pi(q_i) = \Pr[q(t_0) = q_i]. \quad (1.37)$$

1.8.2 | Inference

A particular realisation of a HMM yields a sequence of observations $O = [o(t_1), o(t_2), \dots, o(t_{N_T})]$. The probability of sequence of hidden states $Q = [q(t_0), q(t_1), q(t_2), \dots, q(t_{N_T})]$ given the observed sequence O is denoted $\Pr(Q \mid O)$ and is given by

$$\Pr(Q \mid O) = \Pi[q(t_0)] \prod_{n=1}^{N_T} L[o(t_n), q(t_n); t_n] A[q(t_n), q(t_{n-1}); t_n]. \quad (1.38)$$

A natural question is: what is the most likely sequence of hidden states,

$$Q^* = \max_Q \Pr(Q \mid O)? \quad (1.39)$$

At first blush this seems a daunting task — the space of possible hidden state sequences is exponentially large, $\sim (N_Q)^{N_T}$. However, the Markovian structure allows for an extremely efficient exploration of this space, using the *Viterbi* algorithm (Viterbi, 1967; Rabiner, 1989) — with this approach the overall computational complexity is reduced to $O(N_T N_Q^2)$. The Viterbi algorithm has found widespread use in engineering and scientific applications, from speech recognition (Vargas et al., 2001) to genomics (Antonov and Borodovsky, 2010) to telecommunications (Hayes, 2002). In particular, it has found application in gravitational wave astronomy (Suvorova et al., 2016; Bayley et al., 2019), and in Section 1.9 we discuss the use of HMMs and the Viterbi algorithm in searches for continuous gravitational waves.

A related question which arises in HMM analyses is this: how well does a given model describe a set of observations? That is, what is $\Pr(O | \mathcal{M})$ where \mathcal{M} is a prescribed HMM? As in the case of finding Q^* , the Markovian property makes this apparently difficult question computationally feasible to answer, with the use of the so-called *forward algorithm* (Rabiner, 1989). While the value of $\Pr(O | \mathcal{M})$ is not meaningful in isolation, if we wish to compare models against one another then the ratio $\Pr(O | \mathcal{M}_1) / \Pr(O | \mathcal{M}_2)$, known as the *Bayes factor*, can be used to evaluate the degree to which one model is favoured over another. This approach to model selection is employed in this thesis in conjunction with a HMM-based pulsar timing framework to perform automated pulsar glitch detection (Melatos et al., 2020), and this application is described in more detail in Section 1.10.

Finally, rather than enquiring about the most likely sequence of states, we may wish to know about the probability of any particular state at a particular timestep, given all the data:

$$\gamma_{q_i}(t_n) = \Pr[q(t_n) = q_i | O]. \quad (1.40)$$

Once again an efficient algorithm is known to compute this quantity, known as the *forward-backward* algorithm. As the name suggests, it is an extension of the forward algorithm mentioned above as a means of calculating $\Pr(O | \mathcal{M})$.

1.9 | Hidden Markov models and continuous waves

1.9.1 | The need for hidden Markov models in CW searches

In Section 1.7.2, we introduced the signal model for CW searches as being described by a simple Taylor series, akin to that used in traditional pulsar timing (Equation 1.4). Unfortunately, this is not the whole story — we may well expect that just as timing noise adds an additional stochastic component to the phase evolution in the context of radio pulsar timing, similar phenomena may affect a CW signal from a neutron star. Most directly, we might anticipate the need to incorporate a timing noise component precisely because we observe timing noise in electromagnetic timing observations — if the CW emission is due, for example, to a mountain on the neutron star crust, then the EM and CW emission can be expected to be “locked” to each other, as the EM emission is typically supposed to be tied to the crust by virtue of its high conductivity locking the magnetic field in place. Thus observation of timing noise in EM timing data strongly suggests the possibility of the same phenomenon in the CW signal. Conversely, the CW-emitting component may *not* be locked precisely to the crust, and so even if timing noise

is not observed at a high level in the electromagnetic timing data, wandering may still occur in the CW component.

Given the presence of an extra unknown component in the CW signal phase, one possible approach is to attempt to *track* the underlying signal evolution — adopt a model where some aspect of the unknown signal component (e.g. the frequency) is changing stochastically, and try to estimate its evolution based on the observations and some assumptions about the nature of the stochastic process. This is precisely the kind of problem to which a hidden Markov model is adapted, as long as the underlying stochastic process can be reasonably approximated as Markovian. In the next section we discuss the implementation of HMMs in the CW context in more detail, but first we turn to motivations for including stochasticity in our signal models beyond intrinsic timing noise.

Timing noise as discussed elsewhere in this thesis refers to a process which is not externally driven, but rather arises from some process which is intrinsic to the neutron star. For the purposes of CW searches, the typical magnitude of timing noise observed in non-interacting electromagnetic pulsars (i.e. those not accreting from a binary companion) is relatively small, and may not lead to more than a negligible loss in sensitivity (Jones, 2004). However, *accreting* neutron stars are some of the most promising potential sources for CW emission (see Section 1.7.1), and so it is well worth considering the potential for stochastic evolution in this case also. Although internal timing noise may well be at work in these objects just as it is in non-interacting neutron stars, the external torque due to accretion provides a potentially powerful driver of stochastic behaviour. The structure of the accretion flow is likely to be highly complex and observations of X-ray flux variability as well as direct timing measurements suggest that it is variable on timescales of order days (Bildsten et al., 1997; Mukherjee et al., 2018; Patruno and Watts, 2021; Melatos et al., 2023). Compared to the coherence timescales implied for intrinsic timing noise which are of order months to years — comparable to the length of the data sets being analysed — this much shorter timescale provides a strong motivation for allowing for stochastic wandering in our analyse. It was searches for these accreting systems which motivated the initial development of HMMs in the CW context (Suvorova et al., 2016).

So far the motivation for adopting a HMM-based analysis has been focused on the expectation of a complex noise process in the true underlying signal. However, the flexibility and efficiency afforded by this approach can be useful even when no such process is present. As discussed in Section 1.7.2, in cases where a large region of parameter space is to be covered, a template bank is constructed, containing many templates at different points of parameter space. The density of templates required depends on how quickly the signal-to-noise ratio degrades as the difference between the template and the true signal parameters is increased. By incorporating a “stochastic” component into the signal model and tracking it efficiently, HMM-based approaches can effectively accomodate larger deviations between template and true signal parameters before significant loss of SNR — the “extra” phase contribution due to the parameter mismatch is absorbed into the nominally “stochastic” wandering tracked in the HMM. An extreme example of this is the SOAP pipeline, which tracks the frequency evolution over very short timesteps, typically 30 minutes (Bayley et al., 2019), which allows it to accomodate a very wide range of signal models with very few templates. In particular it is able to directly track the Doppler modulation due to the Earth’s motion, and thus functions as computation-

ally cheap all-sky search (albeit with poorer sensitivity than other pipelines). Another example is a HMM-based search for a signal from the postmerger remnant of the binary neutron star merger GW170817, for which one expects a signal which is so rapidly decreasing in frequency that coherent timesteps of just 1 s are needed (Sun et al., 2019; Abbott et al., 2019a). A less extreme case is presented in Chapter 6, which presents a HMM-based search for unknown neutron stars in five globular clusters. In this application a template bank was required to cover both sky position and spindown rate, and the choice of timestep size was primarily driven by the effect on required template density, not the expected stochastic wandering. Of course, in these cases where managing computational cost is the main motivation, a truly stochastic component can still be accommodated in addition to the contribution from template mismatch. Other semicoherent techniques maintain some sensitivity to stochastic signal behaviour relative to a fully coherent approach even if the stochasticity is not explicitly included in the signal model, but still suffer relative to an HMM-based approach in the face of significant spin wandering (Carlin and Melatos, [submitted](#)).

1.9.2 | Implementing hidden Markov models in the CW context

In almost all HMM-based CW searches, the stochastic component is tracked by allowing the frequency of the signal to wander in a stepwise fashion. In the language set out in Section 1.8.1, we thus make the following identifications:

- The hidden states \mathcal{Q} are the frequencies, evenly spaced within some band Δf .
- The timesteps are evenly spaced over the total timespan to be analysed.
- The set of possible observations \mathcal{O} consists of all possible realisations of the detector data over one timestep, and the sequence of measured observations O is the sequence of detector data streams, broken up according to the timestep boundaries, which are pre-determined.
- The emission likelihood function is some kind of matched filter (e.g. the \mathcal{F} -statistic), computed on each coherent segment, and with all parameters fixed *except* frequency.
- The transition matrix is most often taken as a simple random walk:

$$A(q_i, q_j) = \frac{1}{3} (\delta_{q_{i-1}, q_i} + \delta_{q_i, q_i} + \delta_{q_{i+1}, q_i}). \quad (1.41)$$

- The prior is taken as uniform over all frequencies, $\Pi(q_i) = N_Q^{-1}$.

In Section 1.8.2 we introduced the notion of the most likely sequence of states Q^* (Equation 1.39), which we refer to as a “path”, conditional on the observed data O . Recall that the Viterbi algorithm can be used to efficiently determine Q^* . In practice, we are typically not interested in only the overall most likely path, but the family of most likely paths *conditioned* on the path ending in a given frequency bin, i.e.

$$Q^*(f) = \max_Q \Pr[Q \mid O, q(t_{N_T}) = f]. \quad (1.42)$$

Each path has an associated likelihood $\mathcal{L}(f) = \Pr[Q^*(f) | O]$, and it is these likelihoods which are used to determine the significance of a detection¹⁵.

The framework presented here is the base case where only the frequency is tracked. Extensions to a more complicated HMM structure are possible, e.g. an augmented emission likelihood which includes the possibility of emission at both once and twice the spin frequency of the star (Sun and Melatos, 2019), or an augmented hidden state space and transition matrix which tracks both the frequency and its first time derivative (Sun et al., 2018) or the frequency and phase (Melatos et al., 2021) of the CW signal. The phase-tracking HMM implementation is used in Chapter 6 of this thesis to carry out a search for CW emission from unknown neutron stars in five globular clusters.

1.10 | Hidden Markov models and pulsar timing

1.10.1 | The need for HMMs in pulsar timing

In contrast to the application of HMMs to continuous wave searches, where the aim is detection of a faint signal, in applying HMMs to pulsar timing our aim is wholly to say something about the evolution of the signal. The primary motivation for the HMM pulsar timing framework discussed here is glitch detection (Melatos et al., 2020). Fundamentally this is a question of model selection — is a model with a glitch preferred over a model without? By how much? An important facet of this question is what the “without” model includes. Apart from the effects of barycentring and the spin-down behaviour, which we assume are well-modeled by a standard phase model derived from e.g. TEMPO2, the dominant effect on the pulsar phase is timing noise. In the HMM framework, timing noise is modeled as a dynamical stochastic process which perturbs the spin state of the pulsar, and a “timing noise only” model is a natural base model against which models including a glitch can be compared. Glitches are also naturally interpreted as a change in the dynamical model of the spin state, and thus the whole enterprise of glitch detection in this framework amounts to choosing a model for the spin state evolution. As discussed in the remainder of this section, HMMs provide a convenient computational framework for choosing between models in this way.

1.10.2 | Implementing hidden Markov models for pulsar timing

Rather than tracking only the frequency, as in the typical CW case, for the purposes of pulsar timing we simultaneously track both the frequency and frequency derivative, still allowing them to change in a stepwise manner between timesteps. In this case, again following the discussion in Section 1.8.1, the various HMM components are as follows:

- The timesteps are taken to be the gaps between consecutive barycentred ToA measurements. A generic gap between two ToAs is denoted by z .
- The hidden states \mathcal{Q} are tuples $(\nu, \dot{\nu})$ of pulsar frequencies and frequency derivatives. In particular $(\nu, \dot{\nu})$ specify a *deviation* away from a secular model for the

¹⁵The exact method used to determine significance varies between searches, and we will not cover the range of options here. For the determination of detection significance in the CW search performed as part of this thesis, see Section 6.4.1

evolution of the pulsar spin parameters which is least-squares fit with e.g. TEMPO2. That is, the pulsar frequency evolves as

$$\nu(t) = \nu_{\text{LS}} + \nu + (\dot{\nu}_{\text{LS}} + \dot{\nu})(t - t_0), \quad (1.43)$$

where the LS-subscripted quantities are the values obtained from the least-squares fit and t_0 is a reference epoch.

- The set of possible observations \mathcal{O} is all possible gaps between consecutive ToAs (i.e. the positive reals), and the sequence of measured observations $O = [z_1, z_2, \dots, z_{T_{NT}}]$ in a given realisation is simply the sequence of gaps between the measured ToAs.
- The emission likelihood is a von Mises distribution:

$$\Pr(z \mid \nu, \dot{\nu}) = \frac{\exp\{\kappa \cos[2\pi\phi(z)]\}}{I_0(\kappa)}, \quad (1.44)$$

where $\phi(z)$ is the total phase accumulated (in units of cycles) over the ToA gap of length z , and κ is the “concentration”, a parameter which controls how sharply peaked the distribution is. The distribution peaks when $\phi(t)$ is an integer, i.e. a whole number of turns have elapsed during the gap between the two ToAs constituting the gap.

- The transition matrix is more complex than the CW case, and is derived from an assumed underlying stochastic process driving the evolution of $(\nu, \dot{\nu})$. In the applications presented in this thesis the assumption is that the stochastic evolution is driven by a white noise term in $\ddot{\nu}$, viz.

$$\frac{d^2\nu}{dt^2} = \xi(t), \quad (1.45)$$

$$\langle \xi(t) \rangle = 0, \quad (1.46)$$

$$\langle \xi(t)\xi(t') \rangle = \sigma^2 \delta(t - t'), \quad (1.47)$$

where σ is a parameter which controls the strength of the noise term. A noise process of this kind implies a transition matrix which is a multivariate Gaussian:

$$A(q_i, q_j; t_k) = \frac{1}{\sqrt{2\pi \det \Sigma}} \exp[-(q_i - \mu)^T \Sigma^{-1} (q_i - \mu)], \quad (1.48)$$

where Σ is the covariance matrix and has components

$$\Sigma_{11} = \sigma^2 z_k^3 / 3 \quad (1.49)$$

$$\Sigma_{12} = \Sigma_{21} = \sigma^2 z_k^2 / 2 \quad (1.50)$$

$$\Sigma_{22} = \sigma^2 z_k \quad (1.51)$$

$$(1.52)$$

and μ is the initial state $q_j = (\nu_j, \dot{\nu}_j)$ after propagation to the next timestep under the assumption of pure secular evolution, i.e.

$$\mu = (\nu_\mu, \dot{\nu}_\mu) = (\nu_j + z\dot{\nu}_j, \dot{\nu}_j). \quad (1.53)$$

This is only one possible choice of noise model, used frequently but not exclusively throughout this thesis.

This transition matrix does not include the possibility of a glitch. In a model where a glitch occurs during a particular timestep, the transition matrix is modified to be

$$A(q_i, q_j; t_k) = \frac{1}{N_G} \Theta(\nu_i - \nu_\mu), \quad (1.54)$$

i.e. it is uniform with support for all final states which have frequency larger than the expected frequency under secular evolution – $\Theta(x)$ is the Heaviside step function which is 1 if $\nu_i > \nu_\mu$ and 0 otherwise, and N_G is a normalising constant.

- As in the CW case, the prior is taken to be uniform over the set of hidden states, $\Pi(q_i) = N_Q^{-1}$.

The inference process is somewhat different from the CW case discussed in Section 1.9. Before making any estimate of the evolution of the hidden state, we select between models of the hidden state evolution. Our base model, denoted \mathcal{M}_0 , is the model under which the evolution is driven wholly by the timing noise process, and thus the transition matrix is always given by equation (1.48). This model is then compared against a family of models $\mathcal{M}_1(k)$, wherein the evolution is driven by timing noise *except* in the k th timestep, during which a glitch occurs and the transition matrix is given by equation (1.54). Comparison proceeds via a Bayes factor calculation,

$$K_1(k) = \frac{\Pr[O \mid \mathcal{M}_1(k)]}{\Pr[O \mid \mathcal{M}_0]}, \quad (1.55)$$

where the numerator and denominator are computed using the forward algorithm (see Section 1.8.2). If any of the $K_1(k)$ exceed a pre-defined threshold K_{th} , we say that we have a glitch candidate. Multiple glitches are handled via a greedy procedure – if a glitch candidate is found, we take $k^* = \arg \max_k K_1(k)$ and then compute

$$K_2(k, k^*) = \frac{\Pr[O \mid \mathcal{M}_2(k, k^*)]}{\Pr[O \mid \mathcal{M}_1(k^*)]}, \quad (1.56)$$

where $\mathcal{M}_2(k_1, k_2)$ is the model in which the transition matrix is given by equation (1.54) during the timesteps indexed by both k_1 and k_2 . If again a glitch candidate is found [$K_2(k, k^*) > K_{\text{th}}$ for at least one k], searches for further glitches proceed analogously. The procedure terminates when none of the Bayes factors exceed K_{th} .

Once a model has been selected, we wish to estimate the evolution of the hidden states, in order to be able to extract estimates of the glitch parameters. While we could do this via the Viterbi algorithm, as was done in the CW context, for this application we adopt a different approach and estimate the full posterior distribution over the hidden states at each timestep,

$$\gamma_{q_i}(t_n) = \Pr[q(t_n) = q_i \mid O, \mathcal{M}]. \quad (1.57)$$

This is done using the *forward-backward* algorithm, an extension of the forward algorithm used to compute the Bayes factors in the model selection step (Rabiner, 1989).

1.11 | Thesis outline

The remainder of the thesis is structured as follows.

- Chapter 2 describes a potential pitfall which arises in pulsar timing experiments where the observations of a pulsar are strongly periodic and can lead to misestimation of glitch parameters. We consider the issue in general terms, and show that traditional timing methods are biased towards recovering small glitch sizes, while the HMM-based approach is not biased but may also fail to recover the correct glitch size. We also show that the previously quoted magnitude of a glitch in PSR J1709–4429 is wrong by a factor of approximately 50.
- Chapter 3 presents a search for previously undetected glitches in the first UTMOST data release (Lower et al., 2020), using the HMM glitch detector described in Section 1.10. No new glitches are detected, and systematic upper limits are set on the size of missed glitches for each of the 282 pulsars searched.
- Chapter 4 describes the analysis of 24 years of observations of the Vela pulsar at the Mount Pleasant observatory using the HMM glitch detector, and the discovery of a new small glitch which occurred in January 1992. As in Chapter 3, upper limits are set on the size of missing glitches in this dataset.
- Chapter 5 presents initial results from the pulsar timing programme undertaken as part of the UTMOST-NS project. We give an overview of the programme as a whole, and discuss specific results pertaining to timing noise across the sample, the presence and measurability of second frequency derivatives (i.e. braking indices), and glitches.
- Chapter 6 presents a search for continuous gravitational wave emission from unknown neutron stars in five Galactic globular clusters, using a phase-tracking hidden Markov model and data from LIGO’s third observing run. No significant candidates are found, and we estimate the sensitivity of the search in terms of both strain amplitude and the ellipticity of a putative source. This search represents the first instance of the phase-tracking HMM method being used in a full astrophysical search.
- Chapter 7 summarises the work presented in the thesis and discusses potential avenues for future developments.

2

Effects of periodicity in observation scheduling on parameter estimation of pulsar glitches

This chapter is a reproduction of Dunn et al. (2021a), *Monthly Notices of the Royal Astronomical Society* 504.3, pp. 3399-3411.

Abstract

In certain pulsar timing experiments, where observations are scheduled approximately periodically (e.g. daily), timing models with significantly different frequencies (including but not limited to glitch models with different frequency increments) return near-equivalent timing residuals. The average scheduling aperiodicity divided by the phase error due to time-of-arrival uncertainties is a useful indicator of when the degeneracy is important. Synthetic data are used to explore the effect of this degeneracy systematically. It is found that phase-coherent TEMPO2 or TEMPONEST-based approaches are biased sometimes toward reporting small glitch sizes regardless of the true glitch size. Local estimates of the spin frequency alleviate this bias. A hidden Markov model is free from bias towards small glitches and announces explicitly the existence of multiple glitch solutions but sometimes fails to recover the correct glitch size. Two glitches in the UTMOST public data release are re-assessed, one in PSR J1709–4429 at MJD 58178 and the other in PSR J1452–6036 at MJD 58600. The estimated fractional frequency jump in PSR J1709–4429 is revised upward from $\Delta f/f = (54.6 \pm 1.0) \times 10^{-9}$ to $\Delta f/f = (2432.2 \pm 0.1) \times 10^{-9}$ with the aid of additional data from the Parkes radio telescope. We find that the available UTMOST data for PSR J1452–6036 are consistent with $\Delta f/f = 270 \times 10^{-9} + N/(fT)$ with $N = 0, 1, 2$, where $T \approx 1$ sidereal day is the observation scheduling period. Data from the Parkes radio telescope can be included, and the $N = 0$ case is selected unambiguously with a combined dataset.

2.1 | Introduction

The long-term spin-down of a radio pulsar may occasionally be interrupted by a glitch: an event in which the pulsar’s spin frequency suddenly increases. Glitches are typically recognised by their influence on the timing residuals (Espinoza et al., 2011; Yu et al., 2013), which are the deviations between the expected and measured pulse times of arrival (ToAs). The expected arrival times are predicted by a timing model, which parametrizes both the intrinsic evolution of the rotational phase of the pulsar (frequency and frequency derivatives, glitches) as well as astrometric effects (Roemer, Shapiro, and Einstein delays, dispersion, position, proper motion, and parallax) (Lorimer and Kramer, 2004; Edwards et al., 2006).

To estimate the parameters of a glitch, a model of the effect of the glitch on the rotational phase of the pulsar is assumed. A typical simple form based on a Taylor expansion is (Lower et al., 2020)

$$\Delta\phi_g(t) = \Delta\phi + \Delta f(t - t_g) + \frac{1}{2}\Delta\dot{f}(t - t_g)^2 + \dots, \quad (2.1)$$

where $\Delta\phi_g(t)$ denotes the extra phase accumulated in response to the glitch relative to a no-glitch phase model. The free parameters here are t_g , the glitch epoch; $\Delta\phi$, the permanent jump in rotational phase due to unmodelled effects or an uncertain glitch epoch; Δf , the permanent jump in spin frequency; and $\Delta\dot{f}$, the permanent jump in the first spin frequency derivative with respect to time. These parameters may be estimated in the same way as other parameters in the pulsar timing model: through a least-squares fit which minimizes the χ^2 of the post-fit residuals (Hobbs et al., 2006). Alternatively they may be estimated through Bayesian inference with a software package such as `TEMPONEST` (Lentati et al., 2014), which incorporates parameters describing the deterministic timing model and noise sources and calculates a posterior probability distribution for these parameters via nested sampling with `MULTINEST` (Feroz et al., 2009). When reporting glitch parameters estimated in this way, it is tacitly assumed that the χ^2 of the post-fit residuals has a unique minimum.

In this paper we explore the validity of the single-minimum assumption and the consequences for glitch parameter estimation, when the scheduling of ToA measurements is periodic. By periodic we mean that the gap between consecutive ToAs is approximately equal to an integer multiple of some common period, e.g. if timing data are collected at the same local sidereal time for each observation. The role of observational scheduling on pulsar glitch measurement has received some attention previously. It is well understood that a higher density of observations is advantageous when trying to detect and characterise glitches (Wong et al., 2001; Dodson et al., 2002; Janssen and Stappers, 2006; Ashton et al., 2019a; Basu et al., 2020), but quantitative statements along these lines are rare, due to the large number of factors which may be considered. Espinoza et al. (2014) pointed out that an infrequent observing cadence, combined with a jump in first frequency derivative, can mask a glitch with sufficiently small permanent frequency jumps. They gave a quantitative lower bound on the detectable permanent frequency jump, $\Delta f_{\text{lim}} = \Delta T |\Delta\dot{f}|/2$ where ΔT is the average time between observations. Similarly, Shannon et al. (2016) noted in their analysis of PSR J0835–4510 that glitches with Δf smaller than 10^{-7} Hz are indistinguishable from timing noise, using timing data with

a monthly observing cadence. Melatos et al. (2020) showed that the ability of a hidden Markov model (HMM) to detect glitches with $\Delta f = 10^{-8}$ Hz in the presence of moderate timing noise is diminished, if the time between observations exceeds ~ 10 days. In their timing of the frequently glitching pulsar PSR J0537–6910, Marshall et al. (2004) employed a strategy in which observations are spaced logarithmically, in order to keep phase uncertainty below 0.1 cycles without expending an inordinate amount of observing time. Finally, the Canadian Hydrogen Intensity Mapping Experiment (CHIME) collaboration has recently noted that their determinations of the pulse frequency of newly discovered pulsars may be in error by $n/(1 \text{ sidereal day})$ where n is a small integer, due to the transit nature of the instrument (Good et al., 2021). This is another manifestation of the underlying issue we explore here in detail. To our knowledge, the specific issue of the effect of periodicity in observation scheduling on the estimation of glitch parameters has not been considered previously.

The paper is structured as follows. In Section 2.2 we derive a practical condition for when periodic scheduling leads to errors in estimating the spin frequency of the pulsar, i.e. the time derivative of the phase to leading order in a Taylor expansion. We demonstrate that the condition is satisfied in some existing pulsar timing datasets, leading to a near-degeneracy between timing models with significantly different spin frequencies. In Section 2.3 we specialise to the case of degeneracy between glitch models with different permanent frequency jumps. In Section 2.4 we investigate how degeneracy in timing models affects specific glitch parameter estimation methods. Section 2.4.1 discusses glitch parameter estimation with TEMPO2 (Hobbs et al., 2006), both in the context of phase-coherent timing (Section 2.4.1) and local estimation of the pulsar’s spin frequency (Section 2.4.1). Section 2.4.2 discusses glitch parameter estimation with TEMPO2, and Section 2.4.3 discusses an HMM-based approach (Melatos et al., 2020). Finally, in Section 2.5 we discuss periodic scheduling in the context of the UTMOST public data release (Lower et al., 2020).

2.2 | Phase ambiguity

In this section we explore a simple, non-glitch case in which there may be ambiguity in the measurement of the evolution of the pulsar’s rotational phase. The ambiguity is ultimately due to near-periodicity in the scheduling of the ToA measurements used to determine the timing model parameters. The case without a glitch builds intuition for the more general case with a glitch, which is discussed in Section 2.3.

2.2.1 | Quasiperiodic scheduling

We denote the measured ToAs by $\{t_1, t_2, \dots, t_N\}$, and consider the time gap between the i th and $(i + 1)$ th ToA, $\Delta t_i = t_{i+1} - t_i$. We call a sequence of observations “periodic”, if we have the following approximate equality for all i :

$$\Delta t_i \approx n_i T, \tag{2.2}$$

where n_i is an integer and T is a common, fundamental period independent of i .

Equation (2.2) expresses the periodicity condition intuitively as an approximate equality between the time gap between consecutive ToA measurements, Δt_i , and an integer

multiple of some common period T . However, it is more accurate to think of this condition as a restriction on the fractional part of the factor which multiplies T in the exact version of equation (2.2), which is not strictly an integer. That is, we have the exact equality

$$\Delta t_i = (n_i + \epsilon_i)T, \quad (2.3)$$

where n_i is an integer, and the remainder ϵ_i is a real number satisfying $|\epsilon_i| < 0.5$ by construction.

2.2.2 | Phase error

In the limiting case of perfect periodicity, there is a choice of T for which we have $\epsilon_i = 0$ for all i . The degeneracy is exact: any extra term in the timing model which contributes an integer number of pulsar rotations over the time-scale T predicts ToAs which coincide exactly with those measured, so timing models with and without such an extra term cannot be distinguished. The simplest extra term is of the form

$$\Delta\phi(t) = \frac{Nt}{T}, \quad (2.4)$$

where N is an integer. Equation (2.4) corresponds to a permanent change in the pulse frequency of size $\Delta f = N/T$.

The limiting case $\epsilon_i = 0$ is special. We now allow $\epsilon_i \neq 0$, start with the true phase model $\phi(t)$, and add a spurious term of the form equation (2.4). A term of this form produces indistinguishable timing models for $\epsilon_i = 0$, so one expects the degeneracy to break slightly, when ϵ_i is small. The degeneracy can be quantified by asking: what is the effect of the spurious phase term on the timing residuals, as a function of ϵ_i ? In general the total phase residual at the i th ToA is given by $R_i = \phi_i - n_i$, where ϕ_i is the predicted pulsar rotational phase at the i th ToA, and n_i is the closest integer to ϕ_i (Taylor, 1992). The phase residual contributed across a gap of length $\Delta t_i = (n_i + \epsilon_i)T$ by $\Delta\phi(t)$ is given by

$$\delta\phi_i = \text{frac}[\Delta\phi(\Delta t_i)] = \text{frac}\left[\frac{N}{T}(n_i + \epsilon_i)T\right] = \text{frac}(N\epsilon_i), \quad (2.5)$$

where $\text{frac}(\cdot)$ denotes the fractional part, and the last equality follows from the fact that N and n_i are both integers by construction, so their product has no fractional part. For $N = 1$ we then simply have $\delta\phi_i = \epsilon_i$.

2.2.3 | Worked example

We illustrate the arguments of this section with an example drawn from a real dataset. The UTMOST pulsar timing programme (Bailes et al., 2017) uses data taken at the Molonglo Observatory Synthesis Telescope, and has released public data for a large number of pulsars (Lower et al., 2020). Here we consider the dataset released for PSR J1452–6036, consisting of 287 ToAs measured between July 2017 and July 2019. Our aim is to determine whether the observations of this pulsar are close enough to periodic that there is ambiguity in the phase evolution of the pulsar. We therefore seek to determine the value of T that minimises the average value of $|\epsilon_i|$ as defined in equation (2.3). If there is a




Figure 2.1: Mean fractional scheduling aperiodicity $\langle |\epsilon_i| \rangle$ as a function of scheduling quasiperiod T for the UTMOST observations of PSR J1452–6036.

choice of T which makes $|\epsilon_i|$ particularly small, then the phase residuals induced by a phase term of the form in equation (2.4) will be correspondingly small.

Fig. 2.1 graphs the average of $|\epsilon_i|$ as a function of T for this dataset. A sequence of local minima is clearly visible, with the lowest value of $\langle |\epsilon_i| \rangle$ occurring at $T = 86158$ s, and $\langle |\epsilon_i| \rangle = 4 \times 10^{-3}$. We therefore expect that the magnitude of the timing residuals induced by an extra phase term of the form $\Delta\phi(t) = t/(86158 \text{ s})$ will be $\sim 4 \times 10^{-3}$.

Once we know the magnitude of the residuals induced in a dataset due to a spurious phase term of the form in equation (2.4), we can compare it against the magnitude of the stochastic residuals due to ToA measurement error. The quoted uncertainties on the ToA measurements in the PSR J1452–6036 data are typically 1 ms. The pulsar spins at roughly 6 Hz, so ToA uncertainties ~ 1 ms correspond to phase residuals $\sigma_{\text{ToA}} \sim 6 \times 10^{-3}$. Hence residuals induced by the extra phase term $\Delta\phi(t)$ are of the same order as the residuals due to ToA measurement error. It is therefore reasonable to expect that the two sets of timing residuals look similar, with and without $\Delta\phi(t)$. Equivalently, we refer to the indicative ratio

$$R = \langle |\epsilon_i| \rangle / \langle \sigma_{\text{ToA}} \rangle. \quad (2.6)$$

If the condition $R \lesssim 1$ is satisfied, as it is in this example, we expect phase models with and without the $\Delta\phi(t)$ term to give similar residuals. This is demonstrated in Fig. 2.2, where the top panel shows the timing residuals for the original UTMOST timing model and the middle panel shows the timing residuals for a timing model with an extra phase term $\Delta\phi(t) = t/(86158 \text{ s})$ added. The bottom panel shows the difference between the two sets of residuals. The difference is no more than a few milliseconds, of the same order as the typical ToA error. After re-fitting the spin frequency, the new timing model has a spin frequency 1.1605×10^{-5} Hz larger than the UTMOST timing model. Presented with the top or bottom panel in Fig. 2.2, an analyst would have no reason to doubt either timing model, even though they are different. We emphasise that in this case the degeneracy between the two timing models does not involve glitches.




Figure 2.2: Timing residuals for the UTMOST observations of PSR J1452–6036 between MJD 57955 and MJD 58600, using the UTMOST timing model (*top*), and a model which matches the UTMOST timing model except for an overall frequency increment of 1.1605×10^{-5} Hz (*middle*). The bottom panel shows the difference between the two sets of residuals (error bars not shown). The uncertainty on the spin frequency reported by TEMPO2 is 6×10^{-10} Hz.

2.2.4 | Detecting residuals from frequency misestimation

If the frequency in the timing model is incorrect by a multiple of the inverse of the scheduling period T , the induced phase residual across a given ToA gap is given by equation (2.5). We may therefore seek to detect the “signal vector” \mathbf{S} , the k th element of which is the cumulative sum of the induced phase residuals up to and including the k th ToA gap:

$$S_k = \sum_{i=1}^k \delta\phi_i. \quad (2.7)$$

In the absence of any other sources of timing residuals, \mathbf{S} is exactly what we expect to see if the frequency is misestimated due to periodic observation scheduling. Successful detection of this signal in the residuals enables disambiguation between timing models which produce timing residuals which appear by eye to be indistinguishable.

Our treatment of this detection problem follows Levy (2008). Since we know precisely the form of the expected signal, cross-correlating the vector \mathbf{S} with the residuals $\mathbf{R} = (R_1, R_2, \dots, R_N)$ is an optimal means of detection if the noise in \mathbf{R} is additive, white, and gaussian. We assume that the noise has the covariance matrix $\mathbf{C} = \text{diag}(\sigma_1^2, \sigma_2^2, \dots, \sigma_N^2)$ where σ_i is the reported uncertainty on the i th ToA, measured in cycles. After calculating the expected signal \mathbf{S} according to equation (2.7), we may compute the cross-correlation test statistic

$$\gamma = \mathbf{R}^T \mathbf{C}^{-1} \mathbf{S} \quad (2.8)$$

and compare this test statistic to a threshold γ_{th} chosen to give a particular false alarm rate P_{fa} ,

$$\gamma_{\text{th}} = \sqrt{\mathbf{S}^T \mathbf{C}^{-1} \mathbf{S}} Q^{-1}(P_{\text{fa}}), \quad (2.9)$$

where $Q^{-1}(x)$ is the inverse of

$$Q(x) = \frac{1}{\sqrt{2\pi}} \int_x^{\infty} du \exp(-u^2/2). \quad (2.10)$$

The probability of detection P_d for a given probability of false alarm P_{fa} is then given by

$$P_d = 1 - Q\left[\sqrt{\mathbf{S}^T \mathbf{C}^{-1} \mathbf{S}} - Q^{-1}(P_{\text{fa}})\right]. \quad (2.11)$$

This gives a quantitative estimate of the reliability of detecting the residuals caused by misestimating the frequency when observations are periodically scheduled.

For concreteness, consider again the UTMOST dataset for PSR J1452–6036. We calculate ϵ_i according to (2.3) assuming $T = 86158$ s, and subsequently calculate \mathbf{S} assuming $N = 1$. The latter quantity is the induced phase error signal due to the frequency in the timing model being $1/(86158 \text{ s}) = 1.161 \times 10^{-5}$ Hz too large. For P_{fa} fixed at 0.01, we find $P_d = 0.98$. Hence it is possible to reliably detect the presence of the signal \mathbf{S} due to an incorrect frequency measurement and therefore reject a timing model which contains such a signal, if one is aware of the effect and has a sufficiently long stretch of data. We caution that this is an idealised treatment which ignores the presence of other noise sources (e.g. timing noise). If significant timing noise is present in the data, the assumption of a diagonal noise covariance matrix breaks down, and a modified approach is recommended. We note also that for all of the results presented in this section, N can be either positive or negative – it makes no difference whether the spurious phase term $\Delta\phi(t)$ corresponds to an increase or decrease in frequency.

2.3 | Glitches

We now turn to the subject of degeneracy between glitch models. The idea is essentially the same: under the conditions described in Section 2.2, certain glitch sizes are difficult to distinguish from one another on the basis of ToA measurements, especially based on visual inspection of the timing residuals. We begin by guiding the reader step-by-step through a worked example based on simulated data in Section 2.3.1 in order to illustrate the application of the key ideas. We then show in Section 2.3.2 that a jump in the frequency derivative during a glitch does not affect the argument in Section 2.2 nor the conclusions in the rest of the paper.

2.3.1 | Worked example

We examine a simulated dataset that is generated as follows. First, a timing model is chosen. We choose to take as our starting point the timing model from the UTMOST data release for PSR J1452–6036, in which a glitch was detected at MJD 58600.29 with $\Delta f = 1.745 \times 10^{-6}$ Hz, $\Delta\phi = 0$, and no reported $\Delta\dot{f}$ or exponentially recovering component (Lower et al., 2020). The timing model used to generate the simulated dataset matches the timing model for PSR J1452–6036 from the UTMOST data release, except that the glitch at MJD 58600.29 has size $\Delta f = 1.2106 \times 10^{-5}$ Hz instead. All other glitch parameters are identical – there is no phase jump at the glitch epoch, change in frequency derivative, or exponential recovery included in the simulated data. We wish to generate a set of synthetic ToAs which are consistent with the chosen timing model. Starting with the real PSR J1452–6036 dataset used in Section 2.2, we use `LIBSTEMPO`¹ to generate a new set of idealised ToAs. The synthetic ToAs begin as an exact copy of the PSR J1452–6036 ToAs but are shifted slightly so that they show zero residuals when analysed with the chosen timing model. White, gaussian noise is then added to each ToA at a level commensurate with the reported uncertainty for that ToA. This synthetic dataset serves two purposes. In the remainder of this section it is used to illustrate the principle that periodicity in observation scheduling leads to a degeneracy between glitch models. In sections 2.4.1, 2.4.2, and 2.4.3 it is used to examine how glitch analyses with `TEMPO2`, `TEMPONEST`, or an HMM-based approach may be confounded by this degeneracy between glitch models.

Fig. 2.3 compares two glitch models for the synthetic dataset described above. The top panel of Fig. 2.3 shows the residuals for the true timing model (i.e. the chosen timing model used in generating the synthetic dataset), while the bottom panel shows the residuals for a timing model in which the glitch size is set to 5×10^{-7} Hz. Why 5×10^{-7} Hz? Because the true glitch size of 1.2107×10^{-5} Hz may be written as $1/(86158 \text{ s}) + 5 \times 10^{-7}$ Hz, and hence expresses the degeneracy noted in Section 2.2. The two sets of timing residuals both appear white by eye. The root mean square (rms) residuals for the original timing model with glitch size $\Delta f = 1.2107 \times 10^{-5}$ Hz are $2729 \mu\text{s}$, while the rms residuals for the case with a glitch size of $\Delta f = 5 \times 10^{-7}$ Hz are $2726 \mu\text{s}$. That is, the residuals are nearly equal, even though Δf is approximately 24 times larger in the former model. Note that in the glitch model with $\Delta f = 5 \times 10^{-7}$ Hz, an unphysical phase jump of $\Delta\phi = 0.326$ has also been included.

¹<https://vallis.github.io/libstempo/>





Figure 2.3: Timing residuals for the synthetic dataset described in Section 2.3.1, using a timing model with (*top*) $\Delta f = 1.2106 \times 10^{-5}$ Hz which matches the injected glitch size, and (*bottom*) $\Delta f = 5 \times 10^{-7}$ Hz which is approximately 24 times smaller than the injected glitch.

2.3.2 | Jumps in frequency derivative

Glitches are often accompanied by a jump in the frequency derivative, which can be a significant fraction of the pre-glitch frequency derivative. In Section 2.3.1 we do not include a jump in frequency derivative, to keep things as simple as possible. We now show that jumps in frequency derivative do not affect the arguments made in Section 2.2: the degeneracy persists even when a frequency derivative jump is present.

To be explicit, we consider two glitch models:

$$\phi_{g,1}(t) = \Delta f(t - t_g) + \frac{1}{2}\Delta\dot{f}(t - t_g)^2 \quad (2.12)$$

$$\phi_{g,2}(t) = \left(\Delta f + \frac{1}{T}\right)(t - t_g) + \frac{1}{2}\Delta\dot{f}(t - t_g)^2 \quad (2.13)$$

where Δf , $\Delta\dot{f}$, and t_g are defined as in equation (2.1) and T is the observing period. While these two glitch models contain a frequency derivative change, they nonetheless differ by t/T (plus a constant, t_g/T), which has the form of equation (2.4). Therefore the arguments of Section 2.2 apply, and these two glitch models show comparable residuals when the observations are periodic with period T .

The above argument extends to glitch models with exponentially decaying terms, or any other terms which may be appropriate, as long as those other terms are common to both glitch models being compared, and the only difference between them is a frequency increment as described in Section 2.2.2.

2.4 | Glitch parameter estimation

In this section we discuss the effects of periodic observation scheduling on specific glitch parameter estimation techniques using `TEMPO2` (Section 2.4.1), `TEMPONEST` (Section 2.4.2), and an HMM (Section 2.4.3).

2.4.1 | Tempo2

Phase-coherent timing

In Section 2.3 and Figure 2.3 we demonstrate that when a glitch occurs in a pulsar, there may be multiple timing models which give similarly small timing residuals. Of course, the illustrative worked example in Section 2.3.1 is fashioned deliberately by exploiting our foreknowledge of the glitch parameters injected into a synthetic data set. Does the same outcome arise “naturally”, when the true glitch parameters are unknown, e.g. in real astronomical data or blind injections in synthetic data? In this section we perform a phase-coherent timing analysis using `TEMPO2` of the same synthetic dataset. This is not a truly blind analysis, but the steps taken in the course of this analysis approximate one typical course of action for the parameter estimation of a previously unknown glitch using astronomical data.

We begin by inspecting the timing residuals for the synthetic dataset used in Section 2.3.1 with a timing model that matches the injected parameters but without a glitch. The residuals are shown in the top panel of Fig. 2.4. The pre-glitch residuals are close to




Figure 2.4: Timing residuals for the synthetic dataset described in Section 2.3.1 before (*left panel*) and after (*middle panel*) adding phase jumps to account for wraps in pulse phase. The right panel shows the residuals obtained after accounting for a glitch with $\Delta f = (5.006 \pm 0.015) \times 10^{-7}$ Hz and $\Delta\phi = 0.323 \pm 0.008$.

white, but there is a clear point at MJD 58600 where the timing model “fails”, and the residuals diverge and begin to wrap. At epochs where the phase residuals wrap around, we add phase jumps by hand to restore phase connection. At this stage we use TEMPO2 to tag each ToA with the number of pulses which have elapsed since the first ToA. We emphasise that this pulse-numbering is model-dependent. With the phase jumps added, we see a clear transition from flat phase residuals to a linear ramp in the residuals in the middle panel of Fig. 2.4 – a clean glitch signature. Using the pulse-numbered ToAs it is straightforward to fit the glitch parameters with TEMPO2. For simplicity we fit only the glitch frequency increment Δf and an unphysical glitch phase jump $\Delta\phi$ to account for uncertainty in the glitch epoch. We set the glitch epoch to be the same as the injected glitch epoch, MJD 58600.29, for ease of comparison with the injected glitch parameters. The TEMPO2 fit returns $\Delta f = (5.008 \pm 0.006) \times 10^{-7}$ Hz, and $\Delta\phi = 0.320 \pm 0.003$. The associated residuals are plotted in the bottom panel of Fig. 2.4. By eye, they appear close to white, with variance consistent with the ToA error bars, which are of order 1 ms. This timing model is a good fit: the reduced χ^2 returned by TEMPO2 is 0.98. It is not, however, an accurate recovery of the injected glitch parameters, which are $\Delta f = 1.2106 \times 10^{-5}$ Hz and $\Delta\phi = 0$.

The above analysis exemplifies a general principle: given a set of ToAs containing a glitch, if the ToAs are consistent with multiple glitch models due to periodicity in the observation schedule, then the glitch model with the smallest Δf is more likely to be recovered by a phase-coherent timing analysis. This occurs because ToAs which are consistent with a small glitch display a relatively gentle linear ramp in the post-glitch residuals. This linear ramp will be picked out by eye, and subsequently used to number the post-glitch pulses. A larger glitch (e.g. 1.2106×10^{-5} Hz compared to 5×10^{-7} Hz) has a different phase model, which assigns a different numbering to the post-glitch pulses. Once the analyst restores phase connection (whether correctly or incorrectly) and numbers the pulses according to this phase connected solution, the range of possible glitch models is restricted.

Local frequency estimation

Rather than constructing a phase-connected timing solution and subsequently estimating glitch parameters, it is also possible to estimate the frequency evolution *locally* by fitting for the frequency using small sets of ToAs closely spaced in time which are derived from sub-integrations of a longer observation. In this case, the inter-ToA spacing is much smaller than T , the fundamental observing period. We stipulate that the ToAs are closely spaced in time to sidestep the issue of distinguishing between phase models with frequencies differing by $1/T$, which arises once ToAs separated by more than T are included in the fit. In order to distinguish between timing models which differ in spin frequency by $1/T$, the time τ between the first and last ToA in each local fit must be long enough that the accumulated phase error exceeds the phase error σ_{ToA} due to the ToA uncertainty:

$$\tau/T \gg \sigma_{\text{ToA}}. \quad (2.14)$$

We consider three applications of local frequency estimation to synthetic data. For each application we generate a synthetic dataset with the same timing model parameters and observation schedule as described in Section 2.3.1 (so the observation period T is still

roughly 86158 s). At each session when a synthetic observation is made, we generate four ToAs spaced closely, with $\tau \ll T$. In the first of the three applications, we take $\tau = 1100$ s,² giving $\tau/T \approx 2\sigma_{\text{ToA}}$. The results of the local spin frequency estimation for this dataset are shown in the top panel of Fig. 2.5. The scatter in the post-glitch frequency estimates is on the order of $10 \mu\text{Hz}$, prohibiting a determination of the post-glitch spin frequency which distinguishes between glitch models with sizes separated by $1/T \approx 1.2 \times 10^{-5}$ Hz. However, if τ is increased by a factor of 5, we instead have the result shown in the middle panel of Fig. 2.5. The scatter in the frequency estimation is much reduced to $\sim 1 \mu\text{Hz}$, and it is clear that a glitch with $\Delta f \approx 1.2 \times 10^{-5}$ Hz is preferred over the value recovered by the phase-coherent analysis in Section 2.4.1 ($\Delta f \approx 5 \times 10^{-7}$ Hz). Finally, we can analyse the same synthetic dataset but use eight ToAs spanning two consecutive observing sessions in each local frequency estimate, so that $\tau \approx T$. The results for this case are shown in the bottom panel of Fig. 2.5. The scatter in post-glitch frequency estimations is again reduced, down to roughly $0.1 \mu\text{Hz}$. However, the local estimates are centred around 5×10^{-7} Hz – the same incorrect estimate recovered by the phase-coherent analysis in Section 2.4.1. In both the phase-coherent case and the $\tau \approx T$ case, the assumption that the increase in frequency due to the glitch contributes less than one full rotation between two observation sessions leads to the incorrect estimate.

2.4.2 | TempoNest

The Bayesian pulsar timing package `TEMPONEST` is sometimes used to estimate glitch parameters (Shannon et al., 2016; Yu and Liu, 2017; Parthasarathy et al., 2020; Lower et al., 2020). If the glitch parameters are not already well-constrained, it is common to first use `TEMPO2` to obtain a set of pulse-numbered ToAs, following the procedure described in Section 2.4.1. Note that it is impossible to run `TEMPONEST` without pulse-numbered ToAs, unless the timing model parameters are already well-constrained. As noted in Section 2.4.1, pulse-numbering restricts the range of viable glitch models. Even if one chooses an uninformative prior for Δf , `TEMPONEST` only assigns high posterior probability to glitch models which are consistent with the pulse-numbering used. As such, glitch sizes estimated with `TEMPONEST` exhibit the same bias towards small glitch sizes as those estimated with a phase-coherent `TEMPO2` analysis.

2.4.3 | Hidden Markov model

Recently Melatos et al. (2020) presented a new approach to pulsar glitch detection that models the rotational evolution of the pulsar with an HMM and selects between models with and without glitches present. Like `TEMPO2`-based approaches, the HMM operates on ToAs. Therefore the arguments in sections 2.2 and 2.3 suggest that one ought to be careful when estimating glitch parameters with an HMM, if the observations are scheduled periodically. It turns out that the HMM estimates Δf more accurately than `TEMPO2` and `TEMPONEST` for the synthetic dataset from Section 2.3.1, as reported in Section 2.4.3. However it is still somewhat prone to the same ambiguities arising from periodic scheduling, as demonstrated in Section 2.4.3 with the aid of specific examples.

²This corresponds to one ToA every 4.5 minutes, which is quite typical for pulsars observed by UT-MOST.

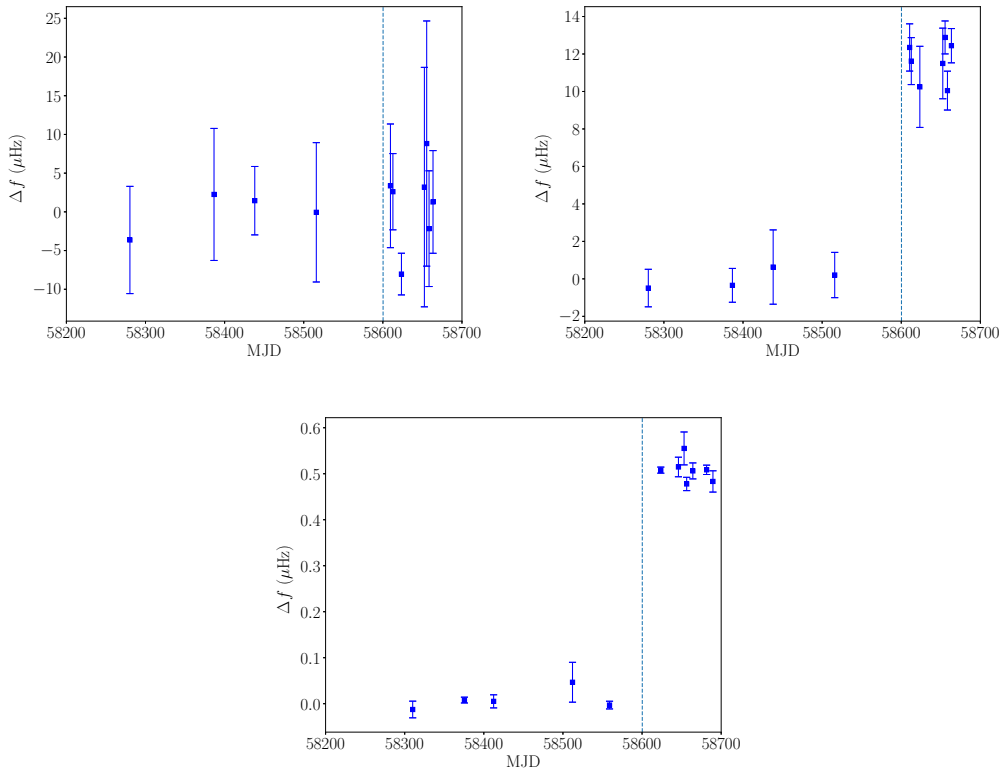


Figure 2.5: Results of local spin frequency estimation described in Section 2.4.1. The top left panel shows the case $\tau/T \approx 2\sigma_{\text{ToA}}$, the top right panel shows the case $\tau/T \approx 10\sigma_{\text{ToA}}$, and the bottom panel shows the case where multiple observing sessions are used in each spin frequency fit. The location of the glitch is indicated by the vertical dotted line. In all cases the true size of the glitch is $\Delta f = 1.21 \times 10^{-5}$ Hz. Note the 10^2 -fold decrease in the vertical scale from the top to the bottom panels.

The HMM consists of three essential components, which we briefly describe here. We refer the reader to Appendix 2.A1 and Melatos et al. (2020) for further details. In the HMM, the state of the pulsar is described by its pulse frequency f , and the time derivative of the pulse frequency \dot{f} . The values of f and \dot{f} are measured relative to a fiducial frequency evolution which is taken from a Taylor expansion of the phase model calculated by TEMPO2. For example, the state $(f, \dot{f}) = (0, 0)$ indicates that the state of the pulsar is exactly as predicted by the Taylor expansion at that timestep. The state (f, \dot{f}) is “hidden”: it is not observed directly, since the HMM operates only on ToAs. Instead, we define an “emission probability” $L(\Delta t_i; f, \dot{f})$ which gives the probability of observing a given ToA gap Δt_i if the pulsar’s state is (f, \dot{f}) during this gap. The expression for $L(\Delta t_i; f, \dot{f})$ used in this paper is given by equation (2.17). Finally, it is necessary to specify the “transition probabilities” which determine the probability that one state transitions to another state after each ToA gap. The HMM tracks the spin wandering directly, as a realization of a Markov chain. This is in contrast to TEMPONEST, which estimates the ensemble characteristics (power spectral density) of the residuals due to timing noise (Lentati et al., 2014). In this work we use transition probabilities which assume a random walk in the second frequency derivative of the pulsar. This matches the prescription adopted by Melatos et al. (2020). Other, qualitatively similar forms of the transition probabilities are also possible and produce qualitatively similar results. Once detected, the parameters of the glitch may be estimated by constructing the sequence of *a posteriori* most likely (f, \dot{f}) states of the pulsar using the forward-backward HMM algorithm, and subsequently reading off Δf at the most probable glitch epoch. With the parameter choices adopted in this work, the HMM is computationally cheap – the search for glitches described in Section 2.4.3 takes roughly 10 min to run on a modern desktop CPU (the quoted time is measured on an Intel Core i5-9300H CPU running at 2.40 GHz.)

Worked example

To illustrate the effect of periodic scheduling on glitch measurement with an HMM, we re-analyse the same synthetic dataset presented in Section 2.3.1 and analysed with TEMPO2 in Section 2.4.1. The details of the parameter choices in the HMM analysis can be found in Appendix 2.A1. Most relevant to the issue at hand is the range of f , which is taken to be $[-3 \times 10^{-7}, 2.5 \times 10^{-5}]$ Hz, bracketing the TEMPO2 timing solution. This range of frequencies is chosen to encompass a large enough range of glitch sizes to demonstrate the essential point. In particular it covers $\Delta f \pm 1/(1 \text{ sidereal day})$, where $\Delta f = 1.21 \times 10^{-5}$ Hz is the injected glitch size. A wider frequency range is possible, but would significantly increase the computation time required without adding anything new [see Section 4.4 of Melatos et al. (2020)]. A glitch is detected at the 231st ToA gap, between MJDs 58599 and 58603, with log Bayes factor 250. We denote this model by $M_1(231)$. While the epoch matches the injected glitch, which occurs at MJD 58600.29, we mimic a realistic analysis by persevering and searching “blindly” for a second glitch by comparing $M_1(231)$ to a set of two-glitch models $M_2(231, k)$ which contain glitches during the 231st gap and the k th gap ($k \neq 231$). This procedure is the greedy hierarchical algorithm described in Section 4.2 in Melatos et al. (2020). We find that the two-glitch model $M_2(231, 232)$ is favored over $M_1(231)$ with a log Bayes factor of 53. We then compare the two-glitch model $M_2(231, 232)$ to a set of three-glitch models $M_3(231, 232, l)$. None of the three-glitch models are favoured over $M_2(231, 232)$ with

log Bayes factor less than zero for all $l \neq 231, 232$, so we terminate the search. Note that we do not interpret $M_2(231, 232)$ as a model with two truly distinct spin-up events, because the ToA gaps are adjacent. $M_2(231, 232)$ is favoured over $M_1(231)$ because the one-glitch models are constructed such that the glitch occurs at the beginning of the ToA gap, whereas in these data the glitch occurs nearly one day into a four-day long gap. The two-glitch model contains enough freedom to mitigate the phase error caused by the simple form of the glitch model assumed by the HMM, which only allows for a glitch to occur at the beginning of a ToA gap, and includes no phase jumps.

Given $M_2(231, 232)$, the HMM forward-backward algorithm computes the posterior distribution of f and \dot{f} during each ToA gap. From this posterior distribution we construct a sequence of frequency and frequency derivative states by choosing the *a posteriori* most likely states during each ToA gap (Rabiner, 1989; Melatos et al., 2020). The frequency sequence $f(t_i)$ constructed in this way is shown in the top panel of Fig. 2.6. The location of the glitch is clear. Reading off the size of the glitch gives $\Delta f = 1.2078 \times 10^{-5}$ Hz, compared to the true glitch size of 1.2107×10^{-5} Hz. While the glitch model allows for a change in \dot{f} (which can be positive or negative), we do not observe such a change in the recovered $\dot{f}(t_i)$ sequence. This is in agreement with the injected glitch parameters, which include $\Delta \dot{f} = 0$.

While this appears to be a relatively successful recovery of the glitch parameters – more so than the phase-coherent analysis of Section 2.4.1 – we can further inspect the posterior distribution of spin states to determine what effect the observation schedule has on this mode of analysis. A heatmap of the logarithm of the posterior distribution of f (marginalised over \dot{f}) as a function of ToA gap index is shown in the bottom panel of Fig. 2.6. It exhibits a multiply peaked structure both before and after the glitch, with peaks separated by 1.1618×10^{-5} Hz. This spacing is significant: it is close to $1/(86158 \text{ s}) = 1.1607 \times 10^{-5}$ Hz, where 86158 s is the observation scheduling period for this dataset. These multiple peaks in the posterior f distribution are indicative of the degeneracy caused by periodic observation scheduling. While the peaks appear to be equal in height based on the logarithmic heatmap, after the glitch the peak at $\Delta f = 1.2078 \times 10^{-5}$ Hz is systematically higher (by a factor of at least 1.5) than the other two, and therefore features in the recovered sequence.

Robustness of parameter estimation against uncertainty in glitch epoch

The HMM succeeds in recovering the correct glitch size in Section 2.4.3 despite the periodic scheduling. However, small variations in the epoch of the glitch can significantly perturb the recovered frequency path when the scheduling is periodic.

Fig. 2.7 shows the marginalised post-glitch posterior distributions of $f(t_{250})$ (t_{250} is chosen to show the posterior distribution of f well after the glitch, which occurs at t_{231}) for 12 injected glitches which have the same size as the previous example, $\Delta f = 1.2107 \times 10^{-5}$ Hz, but epochs distributed uniformly between MJD 58599 and MJD 58603, i.e. anywhere within the 231st ToA gap. The posterior in each case has the same multiply peaked structure as in the bottom panel of Fig. 2.6, with the largest peak located randomly near one of three values: 5×10^{-7} Hz, 1.2×10^{-5} Hz, and 2.4×10^{-5} Hz. In all but one of the panels of Fig. 2.7, only one or two peaks of the three-peak structure seen in the bottom panel of Fig. 2.6 are high enough to be seen by eye.

Thus HMM-based analyses of datasets with periodic scheduling do not necessarily

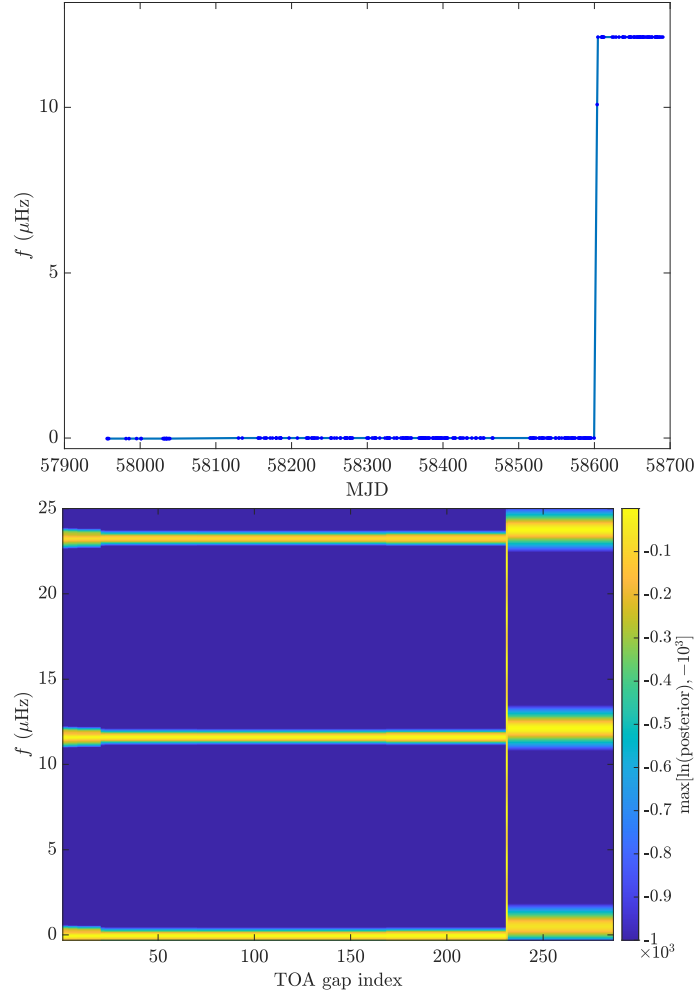


Figure 2.6: (*Top*) Frequency path recovered by the HMM forward-backward algorithm for the analysis of the synthetic dataset described in Section 2.3.1. The recovered glitch size is $\Delta f = 1.2078 \times 10^{-5}$ Hz. (*Bottom*) Heatmap showing the evolution of the logarithm of the posterior probability for f from the HMM analysis of the same synthetic dataset. Time increases from left to right, and at each timestep the posterior has been marginalised over \dot{f} . Values in the heatmap are clipped below -10^3 to aid readability. The favoured model, $M_2(231, 232)$, includes glitches during ToA gaps 231 and 232.

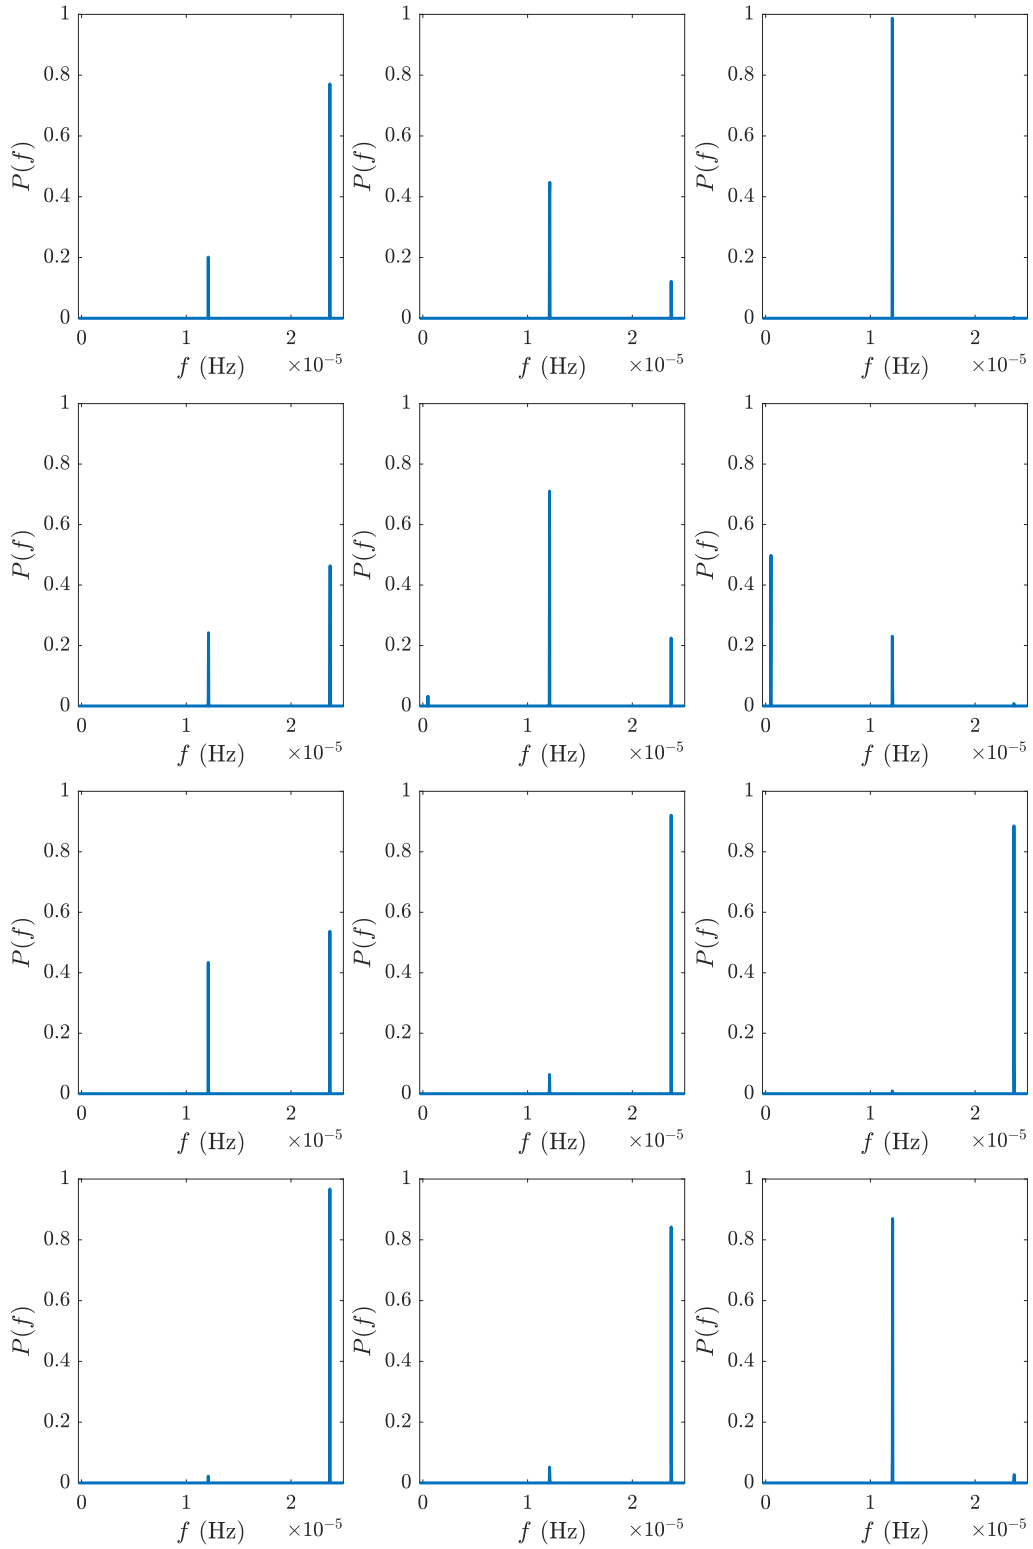


Figure 2.7: Posterior probability distributions $P(f)$ of f during the 250th ToA gap (well after the glitch at t_{231}) for twelve synthetic datasets with randomised glitch epochs in the interval $\text{MJD } 58599 < t_g < \text{MJD } 58603$ as described in Section 2.4.3. In all cases the true frequency deviation is 1.2107×10^{-5} Hz and the glitch occurs during the 231st ToA gap.

recover the correct glitch size: of the twelve posterior f distributions shown in Fig. 2.7, only four peak at the correct location of 1.2107×10^{-5} Hz. The rest have peaks displaced from the true glitch size by approximately $1/(86158 \text{ s})$, indicating that the ambiguity due to periodic scheduling is responsible for the failures to recover the correct glitch size in the other realisations.

Phase-coherent timing analyses of similar datasets are biased towards recovering the smallest plausible glitch size. The HMM shows no systematic bias. Nonetheless it is not guaranteed to return the true glitch size either. Therefore the methods complement one another and are safest to use in tandem. If periodic scheduling is unavoidable for some reason, a chance discrepancy between the methods is one way to catch errors in the estimate of Δf , as the example in Sections 2.4.1, 2.4.1, and 2.4.3 demonstrates in practice. Interestingly the HMM posterior distribution readily reveals the existence of multiple high-likelihood frequency tracks, which are not seen so easily in a timing analysis.

2.5 | Re-assessing UTMOST glitches

From mid-2017 onwards it was decided to use the Molonglo Observatory Synthesis Telescope in transit-only mode (Venkatraman Krishnan et al., 2020). In this mode, the transit of astronomical objects through the primary beam occurs at the same time each sidereal day, with a typical dwell time of 5–20 minutes (Lower et al., 2020). Thus the ToAs used in the UTMOST timing programme are collected in clusters separated by integer multiples of 1 sidereal day (86164 s), closely matching the observing period of 86158 s seen in the UTMOST dataset for PSR J1452–6036. In most cases the cadence of observations is 3–7 days.

With this in mind, we check to see if periodic scheduling compromises the phase reconstruction of UTMOST pulsars. We calculate $\langle |\epsilon_i| \rangle$ [where ϵ_i is defined in equation (2.3)] for each of the 300 pulsars in the UTMOST data release, using only ToAs measured after September 2017 in transit mode, and assuming $T = 86164 \text{ s}$. Recall that $\langle |\epsilon_i| \rangle$ measures the magnitude of timing residuals (measured in terms of cycles) which can be expected if one chooses a timing model which displaces f by $1/T$ from its true value, where T is the scheduling period. Fig. 2.8 shows a histogram of R [defined in equation (2.6)], for the 300 targets. Datasets with $R \gtrsim 1$ are relatively unlikely to support multiple glitch models with comparable residuals. The residuals induced by the wrong choice of Δf are so large, that the error is obvious. However, there is a sizeable population of pulsars with $\langle |\epsilon_i| \rangle / \langle \sigma_{\text{ToA}} \rangle < 1$, including two pulsars which are known to have glitched since September 2017: PSR J1709–4429 ($\langle |\epsilon_i| \rangle / \langle \sigma_{\text{ToA}} \rangle = 0.81$) and PSR J1452–6036 ($\langle |\epsilon_i| \rangle / \langle \sigma_{\text{ToA}} \rangle = 0.12$). The glitch in PSR J1709–4429 has not been reported by any timing programmes other than UTMOST. Given the periodic observation schedule, the possibility that the reported glitch sizes are in error by 1.16×10^{-5} Hz or more must be taken seriously. We now examine the two objects in turn.

2.5.1 | PSR J1709–4429

Lower et al. (2018) reported a glitch in PSR J1709–4429 at $\text{MJD } 58178 \pm 6$ of size $\Delta f / f = (52.4 \pm 0.1) \times 10^{-9}$, based on UTMOST data analysed with TEMPO2 and TEMPOBEST. An updated parameter estimate was published by Lower et al. (2020), who gave a glitch




Figure 2.8: Histogram of R for all 300 pulsars in the UTMOST data release.

size of $\Delta f/f = (54.6 \pm 1.0) \times 10^{-9}$, again based only on UTMOST data. An HMM-based analysis of these data recovers a glitch size of $\Delta f/f = 2405 \times 10^{-9}$, as shown in Fig. 2.9. The glitch size recovered by the HMM is larger than the Lower et al. (2020) result by 2.29×10^{-5} Hz – roughly $2/(1 \text{ sidereal day})$. This suggests that at least one of these methods is confounded by the periodicity of the observations, but without further information it is difficult to decide which glitch size is closer to the truth.

Re-processing the UTMOST data to produce multiple ToAs per observation session allows us to estimate the local spin frequency post-glitch, as described in Section 2.4.1. The results of this exercise are shown in Fig. 2.10. The post-glitch frequency measurements appear to be centred around roughly 2×10^{-5} Hz, consistent with the HMM estimate. Note that preceding the glitch, this pulsar was not observed with sufficiently long observations and sufficient sensitivity to extract enough high-quality ToAs per observing session for a useful local frequency fit. Each of the the pre-glitch frequency estimates therefore incorporate ToAs from observing sessions separated by multiple days, and the error bars are correspondingly much smaller than the post-glitch frequency estimates. The use of widely separated ToAs is not a concern for the pre-glitch frequency estimates, as the pre-glitch frequency is not in question, having been well-measured by UTMOST before the switch to a periodic observing schedule, and this is the first glitch since the switch.

Fortunately, the pulsar timing programme carried out at the Parkes radio telescope has released public data covering the period during which the glitch was reported (Hobbs et al., 2011), and the Parkes pulsar timing programme does not schedule observations with the same regularity as UTMOST. Hence, the combined UTMOST and Parkes data can be expected to estimate the glitch parameters better than the UTMOST data alone. Fig. 2.11 shows timing residuals for two glitch models with the combined UTMOST and Parkes data. With the combined data, the glitch model with $\Delta f/f = 2429.7 \times 10^{-9}$ is clearly preferred, close to what was recovered in the HMM and local frequency estimation analyses.

As a final step, we re-estimate the glitch parameters with the combined UTMOST




Figure 2.9: Recovered frequency path (*top*) and posterior distribution of f (*bottom*) for the HMM analysis of the UTMOST observations of PSR J1709–4429, laid out as in Fig. 2.6. Before the glitch, which occurs during the 81st ToA gap, the posterior distribution of f is well-constrained, showing only a narrow band of support near $f = 0$ Hz. After the glitch, the posterior distribution of f has support in three distinct f regions, separated by $1/(1 \text{ sidereal day})$. The parameters used in this analysis are reported in Table 2.2.




Figure 2.10: Local spin frequency estimation of the glitch size Δf for the glitch in PSR J1709–4429 at MJD 58178 using UTMOST observations, laid out as in Fig. 2.5. The error bars on the pre-glitch frequency estimates are too small to be seen.

and Parkes data using `TEMPONEST` to fit for the timing model parameters and the noise parameters. We start with a pulse numbering derived from a timing model with glitch size $\Delta f/f = 2429.7 \times 10^{-9}$. The glitch model used follows Lower et al. (2020), and is given by

$$\begin{aligned} \Delta\phi_g(t) = & \Delta\phi + \Delta f(t - t_g) \\ & + \frac{1}{2}\Delta\dot{f}(t - t_g)^2 - \Delta f_d\tau_d e^{-(t-t_g)/\tau_d}. \end{aligned} \quad (2.15)$$

The obtained glitch parameters are shown in Table 2.1. The unphysical phase jump $\Delta\phi = -0.342 \pm 0.005$ arises because we choose to fix the glitch epoch at MJD 58178, as in the analysis of Lower et al. (2020). We report only an upper limit on the size of the exponentially decaying term, Δf_d . The posterior distribution for Δf_d has support between the lower end of the prior range, 10^{-18} Hz, and 10^{-9} Hz, but no support above 10^{-9} Hz. The decay timescale τ_d is unconstrained – the posterior distribution has significant support across the entire prior range of (1, 1000) d. Previous glitches of this pulsar have been measured with an exponentially decaying component roughly 1 per cent as large as the permanent frequency jump, with a decay timescale of approximately 100 d (Yu et al., 2013).

2.5.2 | PSR J1452–6036

Lower et al. (2020) also reported a glitch in PSR J1452–6036 at MJD 58600.29 ± 0.05 , with a glitch size of $\Delta f/f = (270.7^{+0.3}_{-0.4}) \times 10^{-9}$ and no measured $\Delta\dot{f}$. In the following we keep $\Delta\dot{f}$ fixed at zero for simplicity.

Local frequency estimation with the UTMOST data does not constrain the glitch size – the typical ToA uncertainty is high, roughly 1 ms, and so the corresponding phase error is $\sim 6 \times 10^{-3}$. Based on the arguments of Section 2.4.1, an observation session long





Figure 2.11: Timing residuals for combined UTMOST and Parkes observations of PSR J1709–4429 for two timing models with glitch sizes $\Delta f/f = 54.6 \times 10^{-9}$ (*top*) and $\Delta f/f = 2429.7 \times 10^{-9}$ (*bottom*). In each plot the unphysical phase jump $\Delta\phi$ has also been adjusted to minimise the jump in residuals before and after the glitch. UTMOST observations are shown as squares, and the Parkes observations are shown as circles. In both plots, all other timing model parameters besides Δf and $\Delta\phi$ are identical to those published in the UTMOST data release.

Table 2.1: Estimated glitch parameters for the glitch in PSR J1709–4429 at MJD 58178. The glitch model parameters are defined in equation (2.15).

| | t_g MJD | $\Delta\phi$ | $(\Delta f + \Delta f_d)/f$ $\times 10^{-9}$ | $\Delta\dot{f}/\dot{f}$ $\times 10^{-3}$ | $\Delta f_d/f$ $\times 10^{-9}$ | τ_d d |
|---------------------|---------------|--------------------|---|---|------------------------------------|-----------------------|
| This work | 58178 ± 6 | -0.342 ± 0.005 | 2432.2 ± 0.1 | 4.7 ± 0.3 | < 0.1 | – |
| Lower et al. (2020) | 58178 ± 6 | 0.372 | 54.6 ± 1.0 | $1.06^{+0.36}_{-0.43}$ | 54.3 ± 1.0 | $99.1^{+11.3}_{-9.6}$ |

enough to break the degeneracy between glitch models would be roughly 1 hr. This is much longer than the actual ~ 5 min observation sessions. The maximum possible observation time in a single transit with UTMOST for this pulsar is somewhat longer, roughly 20 minutes, but not long enough to break the degeneracy. In general the maximum time that UTMOST can observe a given pulsar depends strongly on its declination: pulsars near the ecliptic transit the primary beam in approximately 10 minutes, while pulsars near the south celestial pole remain in the primary beam for hours. However, most pulsars are not routinely observed by UTMOST for more than 10 minutes at a time (Jankowski et al., 2019). Alternatively, the per-ToA uncertainty which would allow a 5 min observation session to break the degeneracy is roughly 0.1 ms.

We may apply the arguments of Section 2.2.4 to search for excess post-glitch phase residuals due to misestimation of the glitch size. We set $P_{\text{fa}} = 0.01$, and find that the threshold calculated according to equation (2.9) is $\gamma_{\text{th}} = 1.99$. The probability of detection given by equation (2.11) is only $P_{\text{d}} = 0.07$. We search for the induced residual signal defined by equation (2.7) in the residuals of three glitch models: the originally reported model, a model with Δf increased by $1/(1 \text{ sidereal day})$ to give $\Delta f/f = 2069 \times 10^{-9}$, and a model with Δf increased by $2/(1 \text{ sidereal day})$ to give $\Delta f/f = 3868 \times 10^{-9}$. The residuals for each glitch model are shown in Fig. 2.12. We find in each case that the test statistic $\gamma = \mathbf{R}^T \mathbf{C}^{-1} \mathbf{S}$ does not exceed γ_{th} , i.e. we do not detect an induced phase error. γ is largest for the glitch model with $\Delta f/f = 3868 \times 10^{-9}$, where we calculate $\gamma = 0.62$. As P_{d} is low, this non-detection does not allow us to constrain the true parameters of the glitch.

Jankowski et al. (2021) recently reported on wideband observations of this pulsar at the Parkes radio telescope which by chance happen to lie on either side of this glitch. Combining these observations with the UTMOST data, they report a glitch size $\Delta f/f = 270.52(3) \times 10^{-9}$, consistent with the value found by Lower et al. (2020). As with the PSR J1709–4429 observations, the Parkes observations of PSR J1452–6036 are not on the same schedule as the UTMOST observations, and so the estimate of Δf from the combined data is not confounded by the periodicity of the UTMOST observations.

2.6 | Conclusion

In this paper we show that periodic scheduling of pulsar observations can lead to erroneous estimates of frequency in pulsar timing models if the frequency is not well-measured *a priori*. We examine in detail the effect this has on the estimation of pulsar glitch parameters. Specifically, the estimated permanent frequency jump Δf may be displaced from its true value by an integer multiple of $1/T$ where T is the scheduling period. We find that in certain existing datasets the excess timing residuals induced by misestimating the spin frequency of the pulsar due to periodic scheduling are comparable to the stochastic residuals induced by ToA measurement error.

We find that “by-eye” attempts to restore phase connection through the use of pulse numbering can bias the recovered glitch size towards smaller values. When the true value of Δf is larger than $1/T$, this bias yields incorrect estimates of Δf . Local frequency estimation can mitigate this bias, as long as the ToAs used in the fits are sufficiently accurate. An HMM-based approach may also fail to recover the correct Δf . However, the HMM does not appear to be biased in the same way as the phase-coherent timing-




Figure 2.12: Timing residuals for UTMOST observations of PSR J1452–6036. Glitches are indicated by the vertical dotted line, with sizes $\Delta f/f = 270 \times 10^{-9}$ (left), 2069×10^{-9} (middle), and 3868×10^{-9} (right). The post-glitch residuals appear to be identical, despite the significant differences in glitch sizes in the three timing models. Other timing model parameters match those released in the UTMOST data release.

based method, and the existence of multiple solutions is readily apparent from a brief inspection of the products of the analysis, e.g. the posteriors of $f(t_1), \dots, f(t_N)$.

We re-evaluate two glitches detected by UTMOST, in PSR J1709–4429 and PSR J1452–6036. In the case of PSR J1709–4429 additional public data from the Parkes radio telescope breaks the degeneracy in glitch models, and we recover a new glitch size $\Delta f/f = (2432.2 \pm 0.1) \times 10^{-9}$, much larger than the previously reported value $\Delta f/f = (54.6 \pm 1.0) \times 10^{-9}$. For PSR J1452–6036, we attempt to detect the presence of phase residuals induced by an incorrect choice of glitch model by cross-correlating the observed residuals with the expected induced signal. However we are unable to definitively confirm or exclude the previously reported glitch model using UTMOST data alone, because the noise is relatively high, and there are relatively few post-glitch observations. Recently released complementary observations from the Parkes radio telescope break the degeneracy and confirm that the glitch size is $\Delta f/f = 270.52(3) \times 10^{-9}$ (Jankowski et al., 2021).

In view of the issues raised here, we recommend that wherever possible, future observing campaigns aimed at glitch measurement should avoid periodic observation scheduling. In cases where periodic observations are unavoidable, such as the CHIME/Pulsar (Ng, 2018; Good et al., 2021) and future UTMOST-2D (Venkatraman Krishnan et al., 2020) programmes, care should be taken when making inferences about the frequency of a pulsar, particularly after a glitch has occurred. A small number of complementary observations from another observatory can also help break the degeneracy between glitch models.

Acknowledgements

Parts of this research are supported by the Australian Research Council (ARC) Centre of Excellence for Gravitational Wave Discovery (OzGrav) (project number CE170100004) and ARC Discovery Project DP170103625. L. Dunn is supported by an Australian Government Research Training Program Scholarship and by the Rowden White Scholarship. M. E. Lower receives support from the ARC Laureate Fellowship FL15010014 and CSIRO Astronomy and Space Science. The authors are grateful to Bill Moran, Rob Evans, and especially Sofia Suvorova for work on the hidden Markov model and useful discussions. The authors are also grateful to the anonymous referee for their helpful comments.

Data availability

The public UTMOST data underlying this work are available at <https://github.com/Molonglo/TimingDataRelease1/>. The other data underlying this work will be shared on reasonable request to the corresponding author.

2.A1 | HMM recipe and parameters

The HMM-based analysis described in Section 2.4.3 involves choosing a number of input parameters. In this appendix we briefly describe these choices. For a more detailed dis-

discussion of the considerations involved in choosing the analysis parameters, see Melatos et al. (2020). There are broadly three classes of parameters involved: those which specify the (f, \dot{f}) pairs under consideration, those which specify the connection between observations (ToAs) and (f, \dot{f}) pairs, and those which specify the probabilities of transitions between (f, \dot{f}) pairs. A complete list of parameters is given in Table 2.2.

We specify the allowed range of frequencies and frequency derivatives in two stages. First, we specify a fiducial phase evolution by fixing the frequency f_0 and frequency derivative \dot{f}_0 at a reference epoch T_0 , derived from a Taylor expansion computed by TEMPO2. The HMM tracks deviations away from this fiducial model on a discrete grid in the f - \dot{f} plane. The range of allowed deviations is specified by lower (f_- , \dot{f}_-) and upper (f_+ , \dot{f}_+) bounds. The discretization is specified by bin sizes η_f and $\eta_{\dot{f}}$.

To incorporate timing noise, we adopt a simple prescription which drives the second frequency derivative with a white noise term $\xi(t)$ satisfying $\langle \xi(t)\xi(t') \rangle = \sigma^2 \delta(t - t')$ [see Section 3.4 in Melatos et al. (2020)]. Other, similar prescriptions yield similar results (Melatos et al., 2020). The free parameter σ controls the magnitude of the timing noise in the model. The discrete nature of the f - \dot{f} grid sets a lower bound on σ : errors in the estimation of \dot{f} , caused by the finite bin width $\eta_{\dot{f}}$, cause the frequency to spuriously wander across a ToA gap by an amount $\delta f_{\text{bin}} \sim \eta_{\dot{f}} \Delta t_i$. It is desirable that the HMM “correct” this spurious wandering, through the freedom allowed in the timing noise model. With $\xi(t)$ defined as above, the frequency wandering across Δt_i is given by $\delta f_{\text{TN}} = \sigma(\Delta t_i)^{3/2}$. Equating δf_{bin} and δf_{TN} with Δt_i replaced by its average over all ToAs, $\langle \Delta t_i \rangle$, suggests the following approximate lower bound,

$$\sigma \geq \eta_{\dot{f}} \langle \Delta t_i \rangle^{-1/2}. \quad (2.16)$$

In the analyses presented here, we choose $\sigma = \eta_{\dot{f}} \langle \Delta t_i \rangle^{-1/2}$.

Given a particular ToA gap Δt_i , the probability of a particular (f, \dot{f}) state is calculated as a von Mises distribution (Melatos et al., 2020)

$$L(\Delta t_i; f, \dot{f}) = [2\pi I_0(\kappa)]^{-1} \exp\{\kappa \cos[2\pi(\Delta t_i f - \Delta t_i^2 \dot{f}/2 + \Delta \Phi_i)]\}, \quad (2.17)$$

where $I_0(x)$ is the modified Bessel function of the first kind, κ is a free parameter known as the concentration, and $\Delta \Phi_i$ is the phase contribution over the ToA gap from the fiducial phase model specified by f_0 , \dot{f}_0 , and the reference epoch T_0 . This distribution is peaked when the number of cycles accumulated across the ToA gap, $\Delta t_i f - \Delta t_i^2 \dot{f}/2 + \Delta \Phi_i$, is an integer (the minus sign appears because we employ a backwards Taylor expansion). To a good approximation we may identify κ as the reciprocal of the squared uncertainty of the phase. There are two important contributions to κ . One is the uncertainty on individual ToAs. For each ToA gap, the phase uncertainties on the two ToAs which bracket the gap, $\sigma_{\text{ToA},1}$ and $\sigma_{\text{ToA},2}$, contribute independently to the total phase uncertainty across the gap. The second contribution comes from the discrete f - \dot{f} grid. The binning in frequency contributes a phase uncertainty $\eta_f \Delta t_i$, and the binning in frequency derivative contributes a phase uncertainty $\eta_{\dot{f}} \Delta t_i^2/2$. Combining all of these uncertainties in quadrature, and recalling the identification of κ with the reciprocal squared phase uncertainty, we arrive at the expression for κ used in these analyses:

$$\kappa = \left[\sigma_{\text{ToA},1}^2 + \sigma_{\text{ToA},2}^2 + (\eta_f \Delta t_i)^2 + (\eta_{\dot{f}} \Delta t_i^2/2)^2 \right]^{-1}. \quad (2.18)$$

Note that κ depends strongly on the length of the ToA gap, and therefore is recalculated for each gap in the dataset. More details on the theoretical underpinnings of (2.17) and (2.18) appear in Section 3.3 and appendix C in Melatos et al. (2020).

Finally, we choose a Bayes factor threshold to be used in model selection via the greedy hierarchical algorithm described in Section 4.2 of Melatos et al. (2020), when models containing glitches are compared against models with no or fewer glitches. Informed by synthetic data tests described by Melatos et al. (2020), we choose here a threshold of $K_{\text{th}} = 10^{1/2}$.

Table 2.2: HMM parameters for the analysis (described in Section 2.4.3) of the synthetic dataset described in Section 2.3.1 and the PSR J1709–4429 analysis described in Section 2.5.1.1. Parameters marked with an asterisk are from the UTMOST data release.

| Parameter | Symbol | Units | Synthetic data | PSR J1709–4429 |
|--------------------------------|---------------------------------------|----------------------|---------------------------------|-------------------------------|
| Timing model reference epoch* | T_0 | MJD | 57600 | 57600 |
| Fiducial frequency* | f_0 | Hz | 6.45193972751 | 9.75429004 |
| Fiducial frequency derivative* | \dot{f}_0 | Hz s^{-1} | -6.03824×10^{-14} | -8.847×10^{-12} |
| Frequency deviation | $[f_-, f_+]$ | Hz | $[-3, 250] \times 10^{-7}$ | $[-3, 250] \times 10^{-7}$ |
| Frequency derivative deviation | $[\dot{f}_-, \dot{f}_+]$ | Hz s^{-1} | $[-6.04, 6.04] \times 10^{-15}$ | $[-4.9, 5.1] \times 10^{-14}$ |
| Frequency bin size | η_f | Hz | 1.7×10^{-8} | 1.7×10^{-8} |
| Frequency derivative bin size | $\eta_{\dot{f}}$ | Hz s^{-1} | 1.21×10^{-15} | 2×10^{-15} |
| Timing noise strength | σ | $\text{Hz s}^{-3/2}$ | 2.56×10^{-18} | 2.05×10^{-18} |
| Mean ToA uncertainty* | $\langle \sigma_{\text{ToA}} \rangle$ | ms | 5.327 | 0.506 |
| Bayes factor threshold | K_{th} | None | $10^{1/2}$ | $10^{1/2}$ |

3

Systematic upper limits on the size of missing pulsar glitches in the first UTMOST open data release

This chapter is a reproduction of Dunn et al. (2022b), *Monthly Notices of the Royal Astronomical Society* 512.1, pp. 1469–1482, reformatted with the following changes:

- A minor typographical error in Section 3.3.1 is corrected.
- Footnote 4 is added in response to a suggestion of the thesis examiner.

Abstract

A systematic, semi-automated search for pulsar glitches in the first UTMOST public data release is presented. The search is carried out using a hidden Markov model which incorporates both glitches and timing noise into the model of the assumed phase evolution of the pulsar. Glitches are detected through Bayesian model selection between models with and without glitches present with minimal human intervention. Nine glitches are detected among seven objects, all of which have been previously reported. No new glitches were detected. Injection studies are used to place 90% frequentist upper limits on the size of undetected glitches in each of the 282 objects searched. The mean upper limit obtained is $\Delta f^{90\%}/f = 1.9 \times 10^{-8}$, with a range of $4.1 \times 10^{-11} \leq \Delta f^{90\%}/f \leq 2.7 \times 10^{-7}$, assuming step events with no post-glitch recoveries. It is demonstrated that including glitch recovery has a mild effect, in most cases increasing the upper limit by a factor of $\lesssim 5$ conservatively assuming complete recovery on a timescale of 100 d.

3.1 | Introduction

The secular electromagnetic spindown of a rotation-powered pulsar is sometimes interrupted by a sudden increase in the spin frequency, known as a glitch. Glitches are often but not always accompanied by a change in the secular spin-down rate and a quasi-exponential recovery (Lyne and Graham-Smith, 2012). The underlying cause of glitches is unknown (Haskell and Melatos, 2015). The standard view of the physical mechanism behind glitches invokes pinning and subsequent unpinning of the vortices of the superfluid neutron component to the lattice of nuclei in the inner crust (Anderson and Itoh, 1975), but this broad picture is by no means certain. Long-term statistical analyses have uncovered interesting features of glitch behaviour both in individual objects (Espinoza et al., 2014; Howitt et al., 2018; Carlin et al., 2019; Ho et al., 2020) and across the pulsar population (Lyne et al., 2000; Melatos et al., 2008; Espinoza et al., 2011; Yu et al., 2013; Fuentes et al., 2017; Melatos et al., 2018). Note that we distinguish between glitches, which involve a jump in the spin frequency, and events involving abrupt changes in frequency derivative associated with magnetospheric changes (Lyne et al., 2010). Although the latter are glitch-like in some respects and interesting in their own right, they are a distinct class of events and we will not search for them explicitly in this work.

Statistical inferences about the glitch phenomenon rely on the completeness of the catalogues of detected glitches. However, the traditional method of glitch detection, which involves identifying a glitch signature “by eye” in a set of timing residuals (Espinoza et al., 2011; Yu et al., 2013), makes it difficult to assess completeness systematically. (Espinoza et al., 2014) and (Yu and Liu, 2017) employed automated glitch detection methods to evaluate detectability limits. The technique presented by Espinoza et al. (2014) has been applied to timing data from the Crab and Vela pulsars (Espinoza et al., 2021), while the technique presented by Yu and Liu (2017) was tested on simulated data but not used to search for glitches in real datasets. Most recently, Singha et al. (2021a) have developed an real-time glitch detection pipeline which operates with minimal human intervention and incorporated it into the timing programme at the Ooty Radio Telescope. Singha et al. (2021a) reported initial tests of detectability limits with this pipeline, and as more data become available a clearer picture will emerge of the completeness of the glitch sample reported as part of this programme.

Melatos et al. (2020) developed a complementary method for pulsar glitch detection which tracks the pulse frequency and frequency derivative with a hidden Markov model (HMM)¹. The HMM selects between models with and without glitches within a Bayesian framework, complementing model selection studies with `TEMPONEST` (Lentati et al., 2014; Shannon et al., 2016; Parthasarathy et al., 2019; Lower et al., 2020). The HMM dynamics include secular spin down and stochastic spin wandering (“timing noise”), as well as step changes associated with glitches (Melatos et al., 2020). As the HMM detects glitches without human intervention, it is well-suited to analysing a large number of pulsar timing datasets. Its speed makes it practical to do injection studies to obtain upper limits on the size of undetected glitches, and hence quantify the completeness of the existing and new catalogues.

In this work we search for glitches in the datasets released as part of the UTMOST pulsar timing programme (Jankowski et al., 2019; Lower et al., 2020). These datasets were

¹https://github.com/ldunn/glitch_hmm

Table 3.1: Observational statistics for the first UTMOST public data release. All quantities are calculated on a per-pulsar basis, and we take the minimum, mean and maximum over the complete set of pulsars in the data release.

| | Minimum | Mean | Maximum |
|------------------------|---------|------|---------|
| Observing timespan (d) | 268 | 1054 | 2024 |
| Cadence (d) | 1.4 | 16 | 49 |
| Number of ToAs | 25 | 107 | 1458 |

released in March 2020, and contain observations of 300 pulsars taken between January 2014 and August 2019. To date, 12 glitches across seven pulsars have been detected as part of this timing programme using traditional methods (Lower et al., 2020). In this paper we search for new glitches beyond those discovered to date and set upper limits on the size of undetected glitches in 282 of the 300 pulsars. The layout of the paper is as follows. In Section 3.2 we briefly describe the data. In Section 3.3 we describe the HMM and explain how to choose the HMM’s control variables and search parameters. In Section 3.4 we present the results of the search. We narrow down the initial list of glitch detections through a veto procedure, and follow up the survivors with a refined HMM analysis to determine the basic glitch parameters: glitch epoch and glitch size. Finally in Section 3.5 we present systematic upper limits on the size of undetected glitches in the UTMOST data release.

3.2 | Data

The UTMOST pulsar timing programme is an ongoing campaign conducted at the Molonglo Observatory Synthesis Telescope, a pair of 778 m long east-west cylindrical paraboloid reflectors located near Canberra, Australia (Bailes et al., 2017). We searched a subset of the data from the first public release² (Lower et al., 2020). The data consist of times of arrival (ToAs) for 300 pulsars, mostly recorded between October 2015 and August 2019, as well as best-fit timing models. We search for glitches only in the 283 pulsars that are not in binary systems, due to difficulties in extracting ToAs from TEMPO2 which are referenced to the reference frame of the pulsar, rather than the solar system barycentre.

The volume and density of available timing data vary significantly between pulsars. Table 3.1 summarises the variation in observing timespan, cadence, and number of ToAs available across the population of UTMOST pulsars. The observing timespan is the time between the first available ToA and the last available ToA for each object, and the cadence is the mean time elapsed between consecutive ToAs. A full description of the observation, data reduction and timing analysis procedures is given by Jankowski et al. (2019).

3.3 | Hidden Markov model

The implementation of a HMM used to search for glitches is described in detail by Melatos et al. (2020). Here we provide a brief description of the most pertinent aspects.

²<https://github.com/Molonglo/TimingDataRelease1/>

Section 3.3.1 discusses the probabilities which determine the dynamics of the HMM, and its connection to the observed data. Section 3.3.2 introduces the Bayesian model selection procedure which is used to select between models with and without glitches present. In Section 3.3.3 we discuss the boundaries and discretisation of the state space of the HMM. The choices for the various parameters are summarised in Table 3.2. These are the choices used in the initial searches for previously unknown glitches (Section 3.4) and in setting upper limits on the sizes of undetected glitches (Section 3.5). Analyses aimed at estimating parameters of detected glitches (Sections 3.4.1–3.4.5) may require different parameter choices on a case-by-case basis, usually in the allowed range of frequencies, which may need to be extended by more than an order of magnitude to accommodate the glitch.

3.3.1 | Transition and emission probabilities

A HMM is an automaton which transitions stochastically between a set of hidden states at discrete times t_1, \dots, t_{N_T} , which are spaced unequally in general. The states are hidden in the sense that they cannot be observed directly; the state of the system must be inferred from observations of auxiliary variables related probabilistically to the hidden states rather than the hidden states themselves. The probability of jumping from state q_i at time t_n to state q_j at t_{n+1} , which is called the transition probability $A_{q_j q_i}(t_n)$, depends only on the state at t_n by the Markov property. The probability that the system occupies the state q_i at time t_n , given an observational datum $o(t_n)$ collected at the same time, is called the emission probability $L_{o(t_n)q_i}$. The prior, Π_{q_i} , is the probability that the system is initialized in the state q_i . Together, $A_{q_j q_i}(t_n)$, $L_{o(t_n)q_i}$, and Π_{q_i} define a HMM uniquely.

To apply a HMM to pulsar timing, we identify the hidden states q_i with a discrete grid of (f, \dot{f}) pairs, which encode the instantaneous spin frequency and its first time derivative. The hidden (f, \dot{f}) states are combined with fixed secular values f_{LS} and \dot{f}_{LS} measured at a reference epoch T_0 to give the instantaneous spin frequency $[f_{\text{LS}} + \dot{f}_{\text{LS}}(t_n - T_0)] + f$ and time derivative $\dot{f}_{\text{LS}} + \dot{f}$ respectively at time t_n . In this work we measure f_{LS} and \dot{f}_{LS} using TEMPO2 (Hobbs et al., 2006). We note that while the HMM requires a measurement of f_{LS} and \dot{f}_{LS} , these values may be derived only from a small subsection of the data in cases where a phase-connected solution spanning the whole dataset is not available. Although phase-connected solutions are available for all of the UTMOST pulsars searched in this work, the HMM does not incorporate the provided pulse numbering information. The states can be enlarged to include the second time derivative \ddot{f} (and higher derivatives), but systematic validation tests with real and synthetic data indicate that state enlargement is unnecessary for the application in this paper (Melatos et al., 2020). The observational datum at each time t_n is the ToA difference $o(t_n) = t_n - t_{n-1}$ [for ease of notation we write $x_n \equiv o(t_n)$ in the remainder of this paper], which is related probabilistically to the hidden states. The prior is deliberately chosen to be flat, i.e. Π_{q_i} is constant within a restricted parameter domain (see Section 3.3.3), as $q_i(t_1)$ is unknown and astrophysically irrelevant. Other structures for the HMM are possible, of course, and the reader interested in pulsar timing is encouraged to experiment with them (Rabiner, 1989).

In the HMM framework, the probability of observing a particular ToA gap x_n if the hidden state of the pulsar is (f, \dot{f}) depends on the accumulated rotational phase over the

Table 3.2: Domain of interest of physical parameters and HMM control parameters.

| Parameter | Symbol | Units | Value |
|--------------------------------|-----------------------|----------------------|--|
| Timing model reference epoch | T_0 | MJD | From UTMOST |
| Secular frequency | f_{LS} | Hz | From UTMOST |
| Secular frequency derivative | f'_{LS} | Hz s^{-1} | From UTMOST |
| Frequency deviation | $[f_-, f_+]$ | Hz | $[-3, 3] \times 10^{-7}$ |
| Frequency derivative deviation | $[f'_-, f'_+]$ | Hz s^{-1} | $[\max(0.1 f'_{\text{LS}}, -10^{-14}), \min(-0.1 f'_{\text{LS}}, 10^{-14})]$ |
| Frequency bin size | η_f | Hz | 4.0×10^{-10} |
| Frequency derivative bin size | η'_f | Hz s^{-1} | $(f'_+ - f'_-)/11$ |
| Timing noise strength | σ | $\text{Hz s}^{-3/2}$ | $\max(10^{-21}, \eta_f \langle x_n \rangle^{-1/2})$ |
| ToA uncertainty | σ_{ToA} | s | From UTMOST |
| Bayes factor threshold | K_{th} | None | $10^{1/2}$ |

gap, $\Phi(t_n; f, f_{\text{LS}}, \dot{f}, \dot{f}_{\text{LS}}, T_0)$ [denoted $\Phi(t_n; \dots)$ for brevity]. Note that $\Phi(t_n; \dots)$ satisfies $0 \leq \Phi(t_n; \dots) \leq 1$ over one period, i.e. it is in units of cycles, not radians. The full expression for $\Phi(t_n; \dots)$ is

$$\Phi(t_n; \dots) = [f + f_{\text{LS}} + \dot{f}_{\text{LS}}(t_n - T_0)]x_n - \frac{1}{2}(\dot{f}_{\text{LS}} + \dot{f})x_n^2. \quad (3.1)$$

The minus sign in the second term arises because we are employing a *backwards* Taylor expansion. Equation (3.1) can be generalised to include a secular second frequency derivative \ddot{f}_0 , if required (Melatos et al., 2020).

If $\Phi(t_n; \dots)$ is close to an integer, the probability of observing the ToA gap x_n is high. This is quantified via a von Mises distribution, in which the probability of observing z given a hidden state (f, \dot{f}) is given by

$$L_{x_n q_i} = \frac{\exp\{\kappa \cos[2\pi\Phi(t_n; \dots)]\}}{2\pi I_0(\kappa)}, \quad (3.2)$$

where $I_0(x)$ is the zeroth modified Bessel function of the first kind, and κ is a parameter known as the concentration. Roughly speaking κ can be thought of as the reciprocal of the variance of $2\pi\Phi(t_n; \dots)$. There are two main contributions to variance in $\Phi(t_n; \dots)$: measurement uncertainty in the ToAs, and the spacing in the discretized f - \dot{f} grid. If the uncertainties in the ToAs at the beginning and end of the gap are $\sigma_{\text{ToA},1}$ and $\sigma_{\text{ToA},2}$ respectively then the contribution to the phase variance is $f_{\text{LS}}^2 (\sigma_{\text{ToA},1}^2 + \sigma_{\text{ToA},2}^2)$. Given spacings in f and \dot{f} of η_f and $\eta_{\dot{f}}$, the respective contributions to the phase variance are $(\eta_f x_n)^2$ and $(\eta_{\dot{f}} x_n^2 / 2)^2$. Combining these contributions in quadrature, we arrive at³

$$\kappa = (2\pi)^{-2} [f_{\text{LS}}^2 (\sigma_{\text{ToA},1}^2 + \sigma_{\text{ToA},2}^2) + (\eta_f x_n)^2 + (\eta_{\dot{f}} x_n^2 / 2)^2]^{-1}. \quad (3.3)$$

During each gap between consecutive ToAs, we assume that the pulsar's hidden state evolves stochastically due to timing noise in the absence of a glitch. The form of the timing noise determines the transition probability $A_{q_j q_i}$ and is unknown a priori for any individual pulsar. One reasonable model, introduced by Melatos et al. (2020) and tested satisfactorily on real data (Melatos et al., 2020; Lower et al., 2021) but certainly not unique, assumes that the timing noise is driven by a white-noise torque derivative,

$$\frac{d^2 f}{dt^2} = \xi(t), \quad (3.4)$$

where $\xi(t)$ is a Langevin term satisfying

$$\langle \xi(t) \xi(t') \rangle = \sigma^2 \delta(t - t') \quad (3.5)$$

and σ is the parameter which controls the strength of the timing noise. Equations (3.4) and (3.5) lead to a robust HMM with easy-to-specify transition probabilities $A_{q_j q_i}(t_n)$. An explicit expression for $A_{q_j q_i}(t_n)$ is given in equations 10–13 and B7–B11 of Melatos

³The factor $(2\pi)^{-2}$ in (3.3) was omitted accidentally by Melatos et al. (2020) in equations (8) and (C3) of the latter reference.

et al. (2020). The choice of (3.4) and (3.5) is pragmatic. Pulsars are not expected to obey (3.4) and (3.5) exactly for many reasons. For example, (3.4) and (3.5) produce Brownian motion in the torque, whereas there is observational evidence that some pulsars exhibit Brownian motion in the frequency, i.e. $df/dt = \xi(t)$ (Cordes and Helfand, 1980; Cordes and Downs, 1985; Parthasarathy et al., 2019).

The choice of σ is also pragmatic and certainly not unique. A detailed study of how to optimize σ on a per-pulsar basis within the HMM framework is beyond the scope of this work; a first pass at some rules of thumb is found in Melatos et al. (2020). Here we follow Melatos et al. (2020) in adopting a simple prescription which is based on the fact that the HMM only tracks \dot{f} to a certain resolution, $\eta_{\dot{f}}$ (see Section 3.3.3). We demand that the necessary ‘‘correction’’ due to binning in the evolution of f across a ToA gap of length x_n is smaller than the dispersion in f caused by the random walk described in (3.4) and (3.5), which implies

$$\sigma = \eta_{\dot{f}} \langle x_n \rangle^{-1/2}, \quad (3.6)$$

where $\langle x_n \rangle$ is the average length of ToA gaps per pulsar. This ensures that the discretisation of \dot{f} does not lead to false alarms, although in some cases it may degrade the performance of the glitch detector by inflating needlessly the strength of timing noise included in the model. Any degradation in sensitivity due to this effect is reflected in the upper limits calculated in Section 3.5. Because the effect of discretisation of \dot{f} is absorbed into the timing noise in the model, we do not expect a significant effect on inferences made on the value of f (e.g the pointwise most likely sequences $\hat{f}(t_n)$ discussed in Section 3.4). The mathematical form of (3.6) is justified in Section 6.1 of Melatos et al. (2020). We additionally follow Melatos et al. (2020) and impose a lower bound $\sigma \geq 10^{-21} \text{ Hz s}^{-3/2}$ to avoid numerical underflow. The prescription described here and listed in Table 3.2 has been successfully tested on synthetic and real data which explicitly does not conform to the timing noise model of equations (3.4) and (3.5) (Melatos et al., 2020; Lower et al., 2021).

If the model includes a glitch during a given ToA gap, then the evolution of the hidden state must be modified accordingly. We adopt the unrestrictive prescription of Melatos et al. (2020): a glitch consists of a positive frequency increment and a possible change in frequency derivative which is allowed to be positive or negative. Explicitly, if the pulsar is in the hidden state (f, \dot{f}) at the beginning of a ToA gap of length x_n , then it is allowed to transition with equal probability to any state (f', \dot{f}') as long as one has $f' > f + \dot{f}x_n$. No restriction is placed on the value of \dot{f}' (within the boundaries specified in Section 3.3.3).

3.3.2 | Model selection

A HMM is a Bayesian inference tool. It works with the fundamental quantity

$$\Pr(Q_{1:N_T} | O_{1:N_T}) = \prod_{q(t_1)} L_{o(t_1)q(t_1)} \prod_{n=2}^{N_T} A_{q(t_n)q(t_{n-1})} L_{o(t_n)q(t_n)}, \quad (3.7)$$

which is the probability that the system occupies the hidden state sequence $Q_{1:N_T} = \{q(t_1), \dots, q(t_{N_T})\}$ given the observation sequence $O_{1:N_T} = \{o(t_1), \dots, o(t_{N_T})\}$ and a model $M = \{A_{q_j q_i}, L_{o(t_n)q_i}, \Pi_{q_i}\}$. The model with no glitch is denoted M_0 , and the model with

a glitch during the k th ToA gap is denoted $M_1(k)$. For a given model M and timing data $D = O_{1:N_T}$ we calculate the model evidence $\Pr(D | M)$ using the HMM forward algorithm (Rabiner, 1989). We can then calculate the ratios

$$K_1(k) = \frac{\Pr[D | M_1(k)]}{\Pr(D | M_0)}, \quad (3.8)$$

for $1 \leq k \leq N_T$, which are Bayes factors, indicating support for each of the N_T glitch-containing models over the no-glitch model. According to Bayes's theorem, the ratio of posterior probabilities of the two models includes an extra factor containing the prior probabilities,

$$\frac{\Pr[M_1(k) | D]}{\Pr(M_0 | D)} = \frac{\Pr[D | M_1(k)]}{\Pr(D | M_0)} \frac{\Pr(M_0)}{\Pr[M_1(k)]}. \quad (3.9)$$

Here we make the simplifying assumption $\Pr[M_1(k)] = \Pr(M_0)$ for all k , and so the Bayes factor $K_1(k)$ and the ratio of posterior probabilities coincide. If $\max_{1 \leq k \leq N_T} K_1(k)$ exceeds a pre-defined threshold K_{th} , we say that we have a glitch candidate. We note briefly that the permissive glitch model used may also accommodate possible abrupt changes in spin-down state (Lyne et al., 2010), and thus model selection may produce glitch candidates associated with these events as well as more typical glitch events.

To account for possible multiple glitches in a dataset, we adopt the greedy hierarchical approach described in Section 4.2 of Melatos et al. (2020). If a glitch candidate is detected when comparing the models $M_1(k)$ to M_0 , we set $k_1^* = \text{argmax}_k K_1(k)$, and then calculate the ratios

$$K_2(k_2) = \frac{\Pr[D | M_2(k_1^*, k_2)]}{\Pr[D | M_1(k_1^*)]}, \quad (3.10)$$

where $M_2(k_1^*, k_2)$ is the model containing two glitches at the k_1^* and k_2 th ToA gaps. If a second glitch candidate is detected, i.e. $\max_{1 \leq k_2 \leq N_T} K_2(k_2) > K_{\text{th}}$, we repeat the procedure, now comparing $M_3(k_1^*, k_2^*, k_3)$ against $M_2(k_1^*, k_2^*)$. The procedure repeats until none of the Bayes factors exceeds K_{th} .

In order to follow up each candidate, we calculate the posterior distribution of frequency and frequency derivative states during each ToA gap using the HMM forward-backward algorithm (Rabiner, 1989). The logic behind this step is discussed in detail in Section 4.3 and Appendix A of Melatos et al. (2020). This posterior distribution can then be used to infer the sequence of most likely frequency states, and hence the most likely size of the frequency jump due to the glitch.

The Bayes factor threshold determines when we have a glitch candidate to be followed up with further analysis. Here we adopt a fixed threshold of $10^{1/2}$, motivated by the synthetic data tests presented in Section 6 of Melatos et al. (2020). $K_{\text{th}} = 10^{1/2}$ gives a false alarm probability of roughly 1%, provided that the timing noise is not much stronger than what is included in the HMM.

3.3.3 | Domain of interest

The domain of interest (DOI) refers to the set of hidden states which are included in the HMM. As in Melatos et al. (2020), we consider only DOIs which form a grid in a "reasonable" $f-\dot{f}$ region⁴, which is restricted to avoid wasteful computation; the state sequence

⁴This is equivalent to setting a uniform prior on f and \dot{f} within a restricted range.

$Q_{1:N_T}$ cannot wander unreasonably far from the TEMPO2 fit f_{LS} and \dot{f}_{LS} , because timing noise and glitches represent modest perturbations on the secular trend. The choices to be made are then the boundaries of the region, and the spacing between points in the grid. The typical DOI parameters used in this work are summarised in Table 3.2.

When searching for unknown glitches and setting upper limits on the size of undetected glitches, the boundary of the f region is chosen to be the same for all pulsars in this study: the region covered is $-3 \leq f / (10^{-7} \text{ Hz}) \leq 3$. This range generously brackets the typical wandering due to timing noise measured in young pulsars to date. The spacing in f is also held fixed for all pulsars at $\eta_f = 4 \times 10^{-10} \text{ Hz}$. This choice represents a trade-off between sensitivity and computational cost. When performing follow-up analysis of a glitch candidate, the range and spacing in f are sometimes modified to encompass the pre- and post-glitch frequencies. In the case of a large glitch, this means increasing the upper boundary of the f region to a value on the order of 10^{-5} Hz . In this case the value of η_f must also be increased to keep the total number of hidden states in the DOI small enough that the computation remains tractable.

The domain of interest and grid spacing in \dot{f} vary between pulsars when searching for new glitches and setting upper limits. Given a measured secular spindown \dot{f}_{LS} , we take the \dot{f} region to be $|\dot{f}| \leq \min(10^{-14} \text{ Hz s}^{-1}, -0.1\dot{f}_{\text{LS}})$. The region is empirically determined, guided by validation experiments with synthetic data, which show that timing-noise-driven excursions in \dot{f} are small compared to \dot{f}_{LS} (Melatos et al., 2020). The \dot{f} spacing $\eta_{\dot{f}}$ is chosen so that there are always 11 points in the domain. This choice is motivated principally by a desire to keep computational cost under control. As a side effect, pulsars with larger values of \dot{f}_{LS} have larger values of σ in the HMM’s timing noise model: a larger \dot{f}_{LS} gives a larger $\eta_{\dot{f}}$, and by equation (3.6) this in turns gives a larger σ . While this relation between \dot{f}_{LS} and σ comports with the astrophysical fact that timing activity is correlated with \dot{f}_{LS} (Arzoumanian et al., 1994; Hobbs et al., 2010; Lower et al., 2020), it does not do so by design. It is a consequence of pragmatic choices which aim to keep false alarms rare [in the case of equation (3.6)] and computational cost low (in the case of the choice of $\eta_{\dot{f}}$).

3.4 | UTMOST glitches

We search for glitches in the UTMOST timing data in three stages. In the initial stage, every pulsar is analysed using the HMM parameters set out in Table 3.2. The results of the initial search are presented in Table 3.3. We report every glitch candidate with a Bayes factor greater than $K_{\text{th}} = 10^{1/2}$. In the second stage, candidates are followed up with a simple veto procedure, and an estimates of the glitch parameters are calculated for those candidates which survive the veto.

Before estimating the parameters of each candidate, we check that the candidate is not due to a transient disturbance. Such a disturbance may be caused by the conditions at the observatory, e.g. a clock error (Verbiest and Shaifullah, 2018). It may also have astrophysical origins. Some pulsars exhibit “mode-changing”, switching between a small number of distinct pulse profiles on timescales of minutes, which can lead to apparent jumps in the pulse phase (Backer, 1970; Helfand et al., 1975; Wang et al., 2007). Changes in propagation through the interstellar medium can also lead to similar apparent phase

Table 3.3: Properties of the initial glitch candidates detected in the first UTMOST data release.

| Object | Epoch MJD | $\ln K_1(k)$ | Vetoed? Y/N |
|------------|----------------|--------------------|----------------|
| J0742–2822 | 57527 ± 2 | 4.05 | Y |
| J0835–4510 | 58218 ± 4 | 4.7×10^4 | N |
| J1105–6107 | 57417 ± 4 | 1.27 | Y |
| J1257–1027 | 58651 ± 10 | 5.4 | N |
| J1359–6038 | 58189 ± 2 | 11.7 | Y |
| J1452–6036 | 58638 ± 1 | 126 | N |
| J1622–4950 | 58076 ± 8 | 8.46 | N |
| J1703–4851 | 58543 ± 21 | 47.5 | N |
| J1709–4429 | 58222 ± 6 | 1.17×10^3 | N |
| J1731–4744 | 58007 ± 2 | 1.20×10^5 | N |
| J1740–3015 | 58393 ± 5 | 4.47×10^4 | N |

jumps (Goncharov et al., 2021). To exclude events of this kind, we re-run the HMM for each candidate using identical parameter choices, but with the ToAs immediately bracketing the candidate removed. If one then obtains $K_1(k) < K_{\text{th}}$, the candidate is vetoed and no further analysis is performed. The results of this veto procedure are noted in the right-most column of Table 3.3. Three candidates are vetoed in this way, leaving eight to be followed up in the second stage of the search. This veto procedure does risk discarding candidates which correspond to true glitches, if the cadence around the candidate is low and the glitch is small. Appendix 3.A1 describes further investigation of each of the three vetoed candidates, in an effort to determine whether they are transient disturbances caused by one of the factors above.

The second stage of the search entails estimating the parameters of each candidate based on the sequence of most likely hidden states (Melatos et al., 2020), which gives the evolution of f , from which the approximate epoch and glitch size can be read off. The follow-up analyses are performed first within an f range which is wider than the range used in the initial analysis. For all seven veto survivors which are ultimately identified as genuine glitch events (i.e. all but the candidate in J1622–4950, see Section 3.4.3), the maximum allowed frequency deviation is 2.5×10^{-5} Hz rather than 3×10^{-7} Hz, so that large glitches are characterised more accurately. The extended f range degrades the sensitivity of the HMM to small glitches, as η_f must increase to keep the computation tractable, which is why a smaller f range is tested in the first stage. After the analysis with an extended f range is complete, each dataset is divided into pre- and post-glitch sections, and these sections are searched again with the smaller f range and η_f used in the initial search. No additional glitch candidates are detected in this way.

The parameter estimation results are summarised in Table 3.4. Two outputs are of particular interest: the sequences of most likely frequency states⁵ $\hat{f}(t_n)$, and the posterior distribution $\gamma_f(t_n)$ of $f(t_n)$. The most likely frequency states are obtained as the modes of the posterior distribution of states q_i at each timestep t_n , where the posterior

⁵The point-wise estimate $\hat{f}(t_n)$ is the most likely value of f at t_n given the data $O_{1:N_T}$. This is subtly different from the f component of the n th element of the most likely sequence of hidden states.

Table 3.4: Properties of the detected glitches confirmed by follow-up analyses. The fractional glitch sizes recovered in the HMM analysis are denoted by $\Delta f/f$, while the values reported by Lower et al. (2020) are denoted by $(\Delta f/f)_{\text{lit.}}$. Phase ambiguity due to periodic observational scheduling prevents inferring $\Delta f/f$ for PSR J1452–6036 (see Section 3.4.2).

| Object | Epoch MJD | $\Delta f/f$ $\times 10^{-9}$ | log Bayes factor | $(\Delta f/f)_{\text{lit.}}$ $\times 10^{-9}$ |
|------------|----------------|----------------------------------|--------------------|--|
| J0835-4510 | 57732 ± 4 | 1436 ± 3 | 7.43×10^4 | $1448_{-0.8}^{+0.9}$ |
| | 58521 ± 7 | 2467 ± 13 | 3.36×10^3 | $2501.2_{-3.2}^{+2.6}$ |
| J1257-1027 | 58650 ± 16 | 2.2 ± 0.4 | 16.4 | $3.20_{-0.57}^{+0.16}$ |
| J1452-6036 | 58606 ± 3 | – | 1.40×10^3 | $270.7_{-0.4}^{+0.3}$ |
| J1703-4851 | 58543 ± 21 | 10 ± 2 | 47.5 | $19.0_{-0.7}^{+1.0}$ |
| J1709-4429 | 58200 ± 27 | 2405 ± 3 | 220 | 54.6 ± 1.0 |
| J1731-4744 | 58007 ± 2 | 3150 ± 14 | 1.40×10^5 | $3149_{-0.4}^{+0.5}$ |
| J1740-3015 | 57476 ± 17 | 225 ± 14 | 1.43×10^3 | $237.7_{-9.3}^{+13.2}$ |
| | 58240 ± 11 | 829 ± 14 | 9.37×10^3 | $842.3_{-5.6}^{+7.1}$ |

is denoted $\gamma_{q_i}(t_n)$ [defined in equation A13 of Melatos et al. (2020)]. The posterior frequency distribution $\gamma_f(t_n)$ is obtained by marginalising $\gamma_{q_i}(t_n)$ over \hat{f} . Plots of $\hat{f}(t_n)$ and $\ln \gamma_f(t_n)$ for the follow-up analyses are shown in Figs. 3.4–3.11.

We can obtain the posterior frequency derivative distribution in much the same way. We do not present these distributions here, as coarse discretisation of \hat{f} in the DOI often leads to unconstraining \hat{f} posteriors. An example is shown in Fig. 3.12, showing $\gamma_f(t_n)$ from the follow-up analysis of J1731–4744 described in Section 3.4. The posterior has support over a significant fraction of the DOI, particularly after the glitch at the 103rd ToA gap, which makes it difficult to draw meaningful conclusions about the evolution of \hat{f} across the glitch.

We remind the reader that the phase model of the HMM is not the same as the phase model of TEMPO2. The HMM allows inter-glitch wandering of f and \hat{f} , and transitions between the hidden f - \hat{f} states occur only at the start of ToA gaps. In addition, the HMM includes no explicit modelling of quasi-exponential post-glitch recovery processes. For all of these reasons, we expect modest discrepancies between the glitch parameters estimated using the HMM and those reported by previous authors, who use TEMPO2 and TEMPONEST to measure the glitch parameters. Those glitches which merit additional discussion are covered in the remainder of this section.

3.4.1 | PSR J0835-4510

PSR J0835-4510 (Vela) exhibits frequent large glitches, at a rate of roughly one every three years with $\Delta f/f \sim 10^{-6}$ typically (Howitt et al., 2018). The UTMOST data available for this pulsar consist of 1420 ToAs recorded between January 9 2014 and September 17 2018. To keep the analysis computationally tractable, we divide the dataset into three sections of approximately 500 ToAs each, with the sections overlapping by 50 ToAs to ensure that any glitches that occur during the gaps between sections were not missed. Details of the section boundaries can be found in Table 3.5. Plots of $\hat{f}(t_n)$ and $\ln \gamma_f(t_n)$ for the

Table 3.5: Data segmentation in the PSR J0835–4510 analysis.

| Section | Start MJD | End MJD |
|---------|-----------|---------|
| 1 | 56666 | 56898 |
| 2 | 56868 | 57606 |
| 3 | 57552 | 58692 |

three sections are shown in Fig. 3.4.

We detected two large glitches in the data, both of which have been reported previously (Palfreyman, 2016; Sarkissian et al., 2019). The first, detected between MJD 57728 and MJD 57734, has $\Delta f/f = (1436 \pm 3) \times 10^{-9}$. The second, detected between MJD 58514 and MJD 58529, has $\Delta f/f = (2467 \pm 13) \times 10^{-9}$. We note a feature which recurs several times throughout these analyses: $\gamma_f(t_n)$ shows multiple peaks for the timesteps following the second glitch, indicating the existence of multiple glitch models which describe the data well, despite being widely separated in frequency. One can see this clearly in the bottom right panel for Fig. 3.4, where the yellow contour splits into three branches for ToA index ≥ 533 . Dunn et al. (2021a) showed that this effect is due to periodicity in the observation schedule. If the separations between consecutive ToAs are nearly integer multiples of a common period T , there can be a degeneracy between glitch models with Δf differing by $1/T$. Since mid-2017 the Molonglo Observatory Synthesis Telescope has operated as a transit instrument (Venkatraman Krishnan et al., 2020), so each pulsar is observed at roughly the same local sidereal time for every observation. Hence, to a good approximation, ToAs recorded by UTMOST after mid-2017 per pulsar are separated by integer numbers of sidereal days. Indeed, the spacing between the peaks in $\gamma_f(t_n)$ in Fig. 3.4 is close to $1/(1 \text{ sidereal day}) = 1.1606 \times 10^{-5}$ Hz.

The degeneracy between glitch models may be alleviated by additional observations which disrupt the periodic scheduling. Fortunately, PSR J0835–4510 is an extremely well-studied pulsar, and the second glitch in the UTMOST dataset has been independently reported by several other facilities (Sarkissian et al., 2019; Kerr, 2019; Gancio et al., 2020). Kerr (2019) estimated the size of the glitch to be $\Delta f/f = (2491.1 \pm 0.5) \times 10^{-9}$ based on data from the *Fermi* Large Area Telescope (Atwood et al., 2009), and Gancio et al. (2020) estimated the size of the glitch to be $\Delta f/f = 2682 \times 10^{-9}$ based on observations taken at the Argentine Institute of Radio Astronomy. These measurements are consistent with the HMM estimate, and with the estimate given by Lower et al. (2020).

3.4.2 | PSR J1452–6036

We detected one glitch in this pulsar, occurring between MJD 58603.6 and MJD 58604.6. The log Bayes factor over the no-glitch model is 1.4×10^3 .

Plots of $\hat{f}(t_n)$ and $\ln \gamma_f(t_n)$ are shown in Fig. 3.6. Inspection of $\gamma_f(t_n)$ indicates that caution is warranted when determining Δf for this glitch. The inferred $\hat{f}(t_n)$ shown in the left panel of Fig. 3.6 suggests a glitch size of $\Delta f/f = 3869 \times 10^{-9}$. However, as with PSR J0835–4510, the three peaks in $\gamma_f(t_n)$ indicate the existence of multiple glitch models which are widely separated in Δf (the separation between peaks is approximately 1.1605×10^{-5} Hz) but nevertheless describe the available data well. This glitch was previously reported by Lower et al. (2020) as occurring at MJD 58600.29(5) with $\Delta f/f =$

$270.7_{-0.4}^{+0.3}$. Using the data in the UTMOST public data release, Dunn et al. (2021a) demonstrated that the available data are consistent with $\Delta f/f = 270 \times 10^{-9} + N/(fT)$ with $N = 0, 1, 2$ and $T \approx 1$ sidereal day, thereby including the Lower et al. (2020) value as one possible option (with $N = 0$). This result is in good agreement with the HMM analysis: the three peaks in the post-glitch frequency posterior generated by the HMM lie at $\Delta f/f = 269 \times 10^{-9}$, 2070×10^{-9} , and 3869×10^{-9} . Fortunately, independent observations at the Parkes radio telescope constrain the size of the glitch well, with Jankowski et al. (2021) measuring $\Delta f/f = 270.52(3) \times 10^{-9}$. This value is consistent with the Lower et al. (2020) estimate and the smallest peak in $\gamma_f(t_n)$. Note that we do not expect the tallest peak in the post-glitch frequency posterior to always correspond to the true glitch size when confounded by periodic scheduling (Dunn et al., 2021a).

3.4.3 | PSR J1622–4950

PSR J1622–4950 is a special object: it is a magnetar which shows large torque variations, with much larger variations in \dot{f} than the DOI $\dot{f}_{\text{LS}} \pm 1 \times 10^{-14} \text{ Hz s}^{-1}$ allowed in the initial search for glitches (Camilo et al., 2018). While this glitch candidate is not vetoed by removing ToAs either side of the glitch, it is probably an artifact caused by the first-pass DOI being too restrictive. We note that the sequences of most likely f and \dot{f} states using the initial search parameters run up against the edges of the DOI. We re-analyse the dataset using a DOI which is somewhat expanded in both f and \dot{f} , with boundaries in f at $f_{\text{LS}} \pm 5 \times 10^{-6} \text{ Hz}$ and boundaries in \dot{f} at $\dot{f}_{\text{LS}} \pm 1 \times 10^{-12} \text{ Hz s}^{-1}$, searching for any glitch candidates in exactly the same way as before. No glitch candidate is detected in this reanalysis, so we do not consider this candidate further. For completeness, Figure 3.7 shows the sequence of most likely frequencies and the posterior frequency probability for the re-analysis with the extended DOI. The posterior has a relatively complex structure, because both the timing noise included in the HMM and the errors on individual TOAs are large, giving the HMM significant freedom in finding viable sequences of hidden states. The large torque variations in PSR J1622–4950 make it difficult (though not impossible) to obtain a phase-connected timing solution covering timespans longer than a few months (Levin et al., 2010). We note briefly that the HMM offers a straightforward method of obtaining the pulse numbering and hence a phase-connected solution: from the sequence of most likely frequencies $\hat{f}(t_n)$ and frequency derivatives $\hat{\dot{f}}(t_n)$ it is easy to calculate the number of pulses during each gap via equation (3.1). From this information the relative pulse numbering is easily derived, and a phase-connected solution obtained.

3.4.4 | PSR J1709–4429

We measured a glitch in this pulsar during the ToA gap between MJD 58172.9 and MJD 58227.7, with size $\Delta f/f = (2405 \pm 3) \times 10^{-9}$. The log Bayes factor over the no-glitch model is 220.

Plots of $\hat{f}(t_n)$ and $\ln \gamma_f(t_n)$ are shown in Fig. 3.9. This detection corresponds to a glitch which was previously reported as occurring at MJD 58178 ± 6 with a glitch size of $\Delta f/f = 54.6 \pm 1.0 \times 10^{-9}$ (Lower et al., 2018; Lower et al., 2020). The glitch reported previously is smaller than the one we recover in this analysis. As with the glitches in PSR J0835–4510 and PSR J1452–6036, the post-glitch frequency posterior is multiply peaked,

with peaks separated by $\sim 1/(1 \text{ sidereal day})$ due to periodic observation scheduling. Dunn et al. (2021a) found that a combined dataset incorporating both the UTMOST data release and observations taken at the Parkes Observatory is consistent with $\Delta f/f = (2432 \pm 0.1) \times 10^{-9}$. Hence we expect that the glitch size recovered by the HMM reflects the true glitch properties.

3.4.5 | PSR J1740–3015

In the follow-up analysis of PSR J1740–3015 we detected two glitches. One glitch occurs in the ToA gap between MJD 57459 and MJD 57486, with size $\Delta f/f = (225 \pm 14) \times 10^{-9}$, and is detected with a log Bayes factor of 1.43×10^3 . The second glitch occurs in the ToA gap between MJD 58229 and MJD 58243, with size $\Delta f/f = (829 \pm 14) \times 10^{-9}$, and is detected with a log Bayes factor of 9.37×10^3 . Plots of $\hat{f}(t_n)$ and $\ln \gamma_f(t_n)$ are shown in Fig. 3.11.

The first glitch was initially reported by Jankowski et al. (2016) based on UTMOST data. They reported $\Delta f/f = (227.29 \pm 0.03) \times 10^{-9}$, consistent with the HMM estimate. Multiple peaks are visible in $\gamma_f(t_n)$ after the second glitch. However, this second glitch was also reported by Basu et al. (2020) based on data taken at the upgraded Giant Metrewave Telescope; they reported $\Delta f/f = 837.4(2) \times 10^{-9}$, consistent with the HMM analysis. No post-glitch recovery term was included in their fit.

Two additional glitches have been reported in this pulsar during the timespan covered by the UTMOST first public release, occurring at MJD 57296.5 ± 0.9 and MJD 57346.0 ± 0.6 with sizes $\Delta f/f = 1.30 \pm 0.04 \times 10^{-9}$ and $\Delta f/f = 1.94 \pm 0.02 \times 10^{-9}$ respectively (Jankowski et al., 2015a; Jankowski et al., 2016). We do not detect these glitches in our analysis. This is not surprising, as the 90% upper limit listed in Table 3.6, namely $\Delta f^{90\%}/f = 41 \times 10^{-9}$ (see Section 3.5), is an order of magnitude larger than the reported sizes of the undetected glitches.

3.5 | Size upper limits

Having detected nine glitches among seven pulsars out of the 283 pulsars searched, we now turn to the question of completeness of this glitch sample. In Section 3.5.1 we discuss the 90% frequentist upper limits set on 282 UTMOST pulsars, which are the main result of this section⁶. In Section 3.5.2 we compare these upper limits to the observed population of glitches, and discuss the completeness of the sample of glitches reported in this work. Finally in Section 3.5.3 we investigate how much quasi-exponential glitch recovery affects the upper limits of Section 3.5.1.

3.5.1 | Frequentist limits

In order to assess completeness, we set 90% frequentist upper limits $\Delta f^{90\%}$ on the sizes of undetected glitches for each pulsar in the UTMOST data release. The upper limit is defined, such that there is a 90% probability of detecting a glitch of size $\Delta f^{90\%}$ in the correct ToA gap, if the glitch occurs at a random epoch distributed uniformly over the

⁶We are unable to set a 90% upper limit for the magnetar PSR J1622–4950, as discussed in Section 3.5.1.




Figure 3.1: Histogram of 90% frequentist upper limits on fractional glitch size for the 282 UTMOST pulsars analysed in this paper.

entire dataset (excluding the first two and last two ToAs). In Sections 3.5.1 and 3.5.2, for the sake of simplicity, we do not include a jump in \dot{f} at the glitch epoch or an exponential post-glitch recovery; the latter effect is considered in Section 3.5.3. The probability of detection for a given pulsar and a given Δf is estimated with 100 synthetic datasets generated using LIBSTEMPO with a randomly chosen glitch epoch and noise injected at the level reported in the UTMOST data release. The injected noise includes both Gaussian ToA measurement error and timing noise, e.g. spin wandering intrinsic to the pulsar (Goncharov et al., 2021). The procedures for generating the synthetic datasets and estimating $\Delta f^{90\%}$ are described in Appendix 3.A2. The prescription for choosing HMM parameters is identical to the prescription for real data, as laid out in Section 3.3. A systematic upper limit injection study of this kind is practical only because the HMM runs fast and automatically without human intervention (Melatos et al., 2020).

Figure 3.1 shows a histogram of $\Delta f^{90\%}/f$ for the 282 pulsars analysed here. The majority (96%) of the $\Delta f^{90\%}/f$ values lie between 10^{-9} and 10^{-7} . The mean fractional upper limit for our sample is $\langle \Delta f^{90\%}/f \rangle = 1.9 \times 10^{-8}$. The minimum $\Delta f^{90\%}/f$ value is 4.5×10^{-11} for the millisecond pulsar PSR J1730–2304, while the maximum is 2.7×10^{-7} in the young pulsar PSR J1123–6259. We do not list every value of $\Delta f^{90\%}/f$ here for readability, but Table 3.6 lists the values for the seven pulsars in which we report at least one glitch in Section 3.4. A complete list of $\Delta f^{90\%}$ values can be found in the Supplementary Materials.

No upper limit is obtained for PSR J1622–4950, where strong timing noise in the HMM (see Section 3.4.3) and a periodic observation schedule mean that it is not possible to attain a detection probability of 90% for any plausible glitch size. Typically, even in the presence of significant timing noise one can increase Δf to a point where timing noise can no longer account for the frequency jump. However, if observations are periodic as they are here, with a period of 1 sidereal day, then glitches with sizes Δf larger than $N/(1 \text{ sidereal day})$ (where N is an integer) are recovered as glitches with size $\Delta f - N/(1 \text{ sidereal day})$. Hence they may never be detected by the HMM if the timing noise

| Object | $\Delta f^{90\%}/f$ $\times 10^{-9}$ |
|------------|---|
| J0835–4510 | 6.7 |
| J1257–1027 | 1.2 |
| J1452–6036 | 14 |
| J1703–4851 | 12 |
| J1709–4429 | 12 |
| J1731–4744 | 62 |
| J1740–3015 | 41 |

Table 3.6: 90% frequentist upper limits on the undetected glitch size in the seven UT-MOST pulsars for which at least one glitch is detected by the HMM.

is large enough to account for a change in frequency of $\Delta f - N/(1 \text{ sidereal day})$.

For the data analysed here, the sensitivity of the HMM is principally controlled by the largest observing gaps in the data. Large gaps allow for significant deviations in the spin frequency to be absorbed into the timing noise model in the HMM, with the expected fractional upper limit proportional to $\dot{f}_+ \max_n x_n$. When this effect dominates, a regular and frequent observing cadence is most useful in obtaining more stringent upper limits. However, we caution that this is not true in all regimes. When the observing cadence is short enough, or the allowed wandering due to timing noise is small enough, the sensitivity of the HMM is instead controlled by the phase uncertainty incorporated into the HMM via the κ parameter (see Section 3.3.1). In this case the expected fractional upper limit is roughly proportional to $2\pi\kappa^{-1/2}\langle x_n \rangle^{-1}$, and increasing the cadence (decreasing $\langle x_n \rangle$) further will give *lower* sensitivity, if κ is dominated by the contribution from σ_{ToA} , which is independent of x_n . We refer the reader to section 6 and appendix G of Melatos et al. (2020) for further discussion on the sensitivity of the HMM glitch detector.

3.5.2 | Population-level comparison

We compare the upper limits obtained in Section 3.5.1 to the observed size distribution aggregated across the entire pulsar population, as recorded in the Jodrell Bank Observatory (JBO) glitch catalogue⁷ (Espinoza et al., 2011). This is not exactly a like-for-like comparison: the catalogue of observed glitches combines a wide variety of datasets and analyses, with varying observation scheduling and glitch detection strategies. Nevertheless it is instructive to ask what categories (if any) of glitches observed in other pulsars are not detectable by the HMM in the UTMOST data release.

Fig. 3.2 shows a histogram of all $\Delta f/f$ values listed in the JBO glitch catalogue overlaid with a histogram of the upper limits obtained in this analysis. There is a population of detected glitches in the JBO catalogue with $10^{-12} \lesssim \Delta f/f \lesssim 10^{-9}$ which are smaller than the 90% upper limits obtained for most of the pulsars in our sample. Glitches in this size range are unlikely to be detected by the search in this paper. Of course, this raises the interesting question of whether some glitches in the JBO catalogue with $\Delta f/f \lesssim 10^{-9}$ are false alarms. This comes down to distinguishing timing noise from glitches through Bayesian model selection and calculating $\Delta f^{90\%}$ for the relevant observational studies

⁷<http://www.jb.man.ac.uk/pulsar/glitches/gTable.html>




Figure 3.2: Histogram of size upper limits $\Delta f^{90\%}/f$ computed in Section 3.5.1 (orange) compared with sizes $\Delta f/f$ detected in the entire pulsar population in the Jodrell Bank Observatory catalogue (blue).

in the literature, a task which is challenging without an unsupervised algorithm like the HMM (Janssen and Stappers, 2006; Chukwude and Urama, 2010; Espinoza et al., 2014; Yu and Liu, 2017).

By way of comparison, we briefly highlight three other investigations of glitch detectability which are similar in spirit to the current work. Janssen and Stappers (2006) performed Monte Carlo injections for a single pulsar, PSR J0358+5413, finding that glitches as small as $\Delta f/f = 10^{-11}$ can be detected by eye. It is unclear, however, how confidently such glitches can be detected. Espinoza et al. (2014) employed an automated glitch detection algorithm to search for glitches in the Crab pulsar, and reported a minimum glitch size in that case, which is intrinsic to the pulsar and not an artifact of the detector performance. This technique has also been applied to the Vela pulsar, and a lack of small glitches was reported in that case also (Espinoza et al., 2021). However, the detector returns a large number of (anti-)glitch candidates which must be classified after the fact as timing noise or glitches. This obstructs the characterisation of the detector in a controlled environment via a suite of synthetic data tests, as well as the application of this technique to a large number of datasets. Yu and Liu (2017) assessed the completeness of the glitch catalogue reported in Yu et al. (2013) through Monte Carlo simulations in which glitch detection was performed using `TEMPONEST`. They concluded that the reported glitch catalogue contains all glitches detectable by manual inspection of timing residuals. However the criterion for a positive detection relies on knowing the true glitch epoch; it cannot be extended to finding previously unknown glitches. We emphasise that although some glitches may be missed, the upper limits calculated here are derived from simulated searches of every pulsar individually.

3.5.3 | Post-glitch recovery

Many (but not all) glitches exhibit a degree of recovery over timescales of days to months, such that part (or all) of Δf reverses, leaving a permanent frequency jump Δf_p (Shemar and Lyne, 1996). The recovery is typically modelled as one or more exponential terms in the post-glitch frequency evolution, viz.

$$f(t) = f(0) + \Delta f_p + \sum_{i=1}^M \Delta f_i e^{-t/\tau_i} \quad (3.11)$$

for a hypothetical glitch occurring at $t = 0$, where the Δf_i are the sizes of the M exponentially recovering components, with recovery timescales τ_i . In many events one has $M = 1$, but where the pulsar is well-observed following the glitch, more exponential terms may be incorporated, e.g. $M \leq 4$ (Dodson et al., 2002). The glitch population as a whole exhibits a wide variety of recovery behaviour. Values of the healing parameter

$$Q = \frac{\sum_i \Delta f_i}{\Delta f_p + \sum_i \Delta f_i} \quad (3.12)$$

are typically between 0 and 1, with $Q \ll 1$ being more common for large glitches ($\Delta f_p/f \gtrsim 10^{-6}$) (Yu et al., 2013).

In Section 3.5.1 we set 90% frequentist upper limits on the sizes of undetected glitches assuming a glitch model (both in the simulated data and the HMM) with no recovery (i.e. $Q = 0$). We now investigate whether including recovery in the simulated data significantly changes these upper limits. To this end, we recompute 90% frequentist upper limits as in Section 3.5.1, but now assume that the glitch recovers completely (i.e. $Q = 1$ in the simulated data) on a timescale of $\tau_1 = 100$ d, typical of many pulsars (Yu et al., 2013). Note that we do not modify the phase model of the HMM in any way – no attempt is made to model the exponential recovery as part of the glitch detection step. Fig. 3.3 shows a histogram of the ratios between the $Q = 1$ and $Q = 0$ values of $\Delta f^{90\%}$ for each pulsar. In the majority of cases the effect of recovery is not severe: for 87% of the pulsars one obtains $\Delta f_{Q=1}^{90\%} < 5\Delta f_{Q=0}^{90\%}$. Hence we do not expect complete glitch recovery on month-long timescales to affect significantly the results presented in Section 3.5.1. The outliers with large ratios $\Delta f_{Q=1}^{90\%}/\Delta f_{Q=0}^{90\%}$ typically have small intrinsic \dot{f}_{LS} , with $\dot{f}_{\text{LS}} > -10^{-15} \text{ Hz s}^{-1}$. In this regime the extra \dot{f} due to quasi-exponential recovery is much larger than the range of \dot{f} in the DOI, which covers only $\pm 0.1 \dot{f}_{\text{LS}}$ (see Section 3.3.3). Hence the HMM struggles to track the evolution of f and \dot{f} , and is correspondingly less sensitive. Conversely, for six pulsars we have $0.8 < \Delta f_{Q=1}^{90\%}/\Delta f_{Q=0}^{90\%} < 1$. In all but one case these pulsars are monitored with relatively high cadence, and have sufficiently wide ranges in the \dot{f} DOI to allow the tracking of the extra \dot{f} from the exponential recovery⁸ Hence we expect that the change in \dot{f} from the recovery allows the HMM to detect glitches more readily in this regime.

⁸In the remaining case, which is PSR J1736–2457, we expect that simple statistical fluctuation is the cause: the observing cadence is relatively low, with several gaps of 10–60 days present in the data. In this case, when relaxations are included and the effect of a glitch on f and \dot{f} can decay away within one or two post-glitch ToAs, the estimated value of $\Delta f^{90\%}$ may depend somewhat on the epochs of the injected glitches in the synthetic datasets.




Figure 3.3: Histogram for 282 UTMOST pulsars of the ratio of 90% glitch size upper limit computed with a completely recovering glitch ($\Delta f_{Q=1}^{90\%}$) to the 90% upper limit computed with no recovery ($\Delta f_{Q=0}^{90\%}$).

The above analysis only considers one point in the space of possible (Q, τ_1) choices (though it is a fairly typical point). A full exploration of the Q - τ_1 plane for every pulsar in the UTMOST data release is beyond the scope of this paper. We emphasise that the conclusions drawn about upper limits here and in Section 3.5.1 are conditional not only on the analysis method but also on the assumed glitch model, a feature of any timing analysis.

3.6 | Conclusion

In this paper we present a search for glitches using a HMM in 283 pulsar timing datasets released by the UTMOST pulsar timing programme, covering observations taken between October 2015 and August 2019. We detect nine glitches among seven pulsars, all of which have been previously reported. The inferred Δf is usually consistent with previous discoveries, except when there is ambiguity due to near-periodic scheduling (Dunn et al., 2021a). In this case the discrepancy in Δf can be large, viz. an integer multiple of T^{-1} , where T is the observation scheduling period. For all the glitches detected in this work, complementary observations by other observatories allow the unambiguous determination of Δf . In principle the Δf inferred from an HMM analysis may be biased by quasi-exponential post-glitch recovery, which is not included in the HMM in its current implementation. However, this effect is demonstrated to be small for the glitches measured in this work, typically $\lesssim 2\%$ for the 7 objects studied here. Incorporating post-glitch recoveries into the HMM (at the expense of introducing new parameters) is a priority for future work.

For each object, we perform injection studies to set frequentist upper limits on the size of undetected glitches. The mean 90% upper limit on the fractional size of undetected glitches is $\langle \Delta f^{90\%} / f \rangle = 1.9 \times 10^{-8}$. The smallest value of $\Delta f^{90\%} / f$ is 4.1×10^{-11} ,

Table 3.7: Top and bottom five $-\dot{f}\langle\Delta t\rangle$ values among pulsars observed by UTMOST. The Pearson correlation coefficients between $\Delta f/f$ and the forward (r_+) and backward (r_-) waiting times are also listed. The objects in the top (bottom) half of the table are more likely to exhibit cross-correlations between glitch size and forward (backward) waiting time.

| Object | N_g | $-\dot{f}\langle\Delta t\rangle$ (Hz) | r_+ | r_- |
|------------|-------|---------------------------------------|-------|-------|
| J1709–4429 | 5 | 1.8×10^{-3} | 0.66 | −0.47 |
| J1803–2137 | 6 | 1.4×10^{-3} | 0.91 | −0.24 |
| J0835–4510 | 20 | 1.2×10^{-3} | 0.24 | 0.55 |
| J1048–5832 | 6 | 6.0×10^{-4} | 0.58 | −0.48 |
| J1105–6107 | 5 | 4.2×10^{-4} | 0.87 | −0.37 |
| J1731–4744 | 5 | 4.4×10^{-5} | −0.71 | 0.99 |
| J1740–3015 | 36 | 3.5×10^{-5} | 0.29 | −0.02 |
| J1705–1906 | 4 | 1.2×10^{-5} | 0.97 | −0.54 |
| J1825–0935 | 7 | 1.1×10^{-5} | 0.91 | −0.30 |
| J1902+0615 | 6 | 2.0×10^{-6} | 0.49 | −0.31 |

calculated for the millisecond pulsar PSR J1730-2304. The largest value of $\Delta f^{90\%}/f$ is 2.7×10^{-7} , calculated for the young pulsar PSR J1123–6259. Obtaining more stringent upper limits using the existing data would require a more complete characterisation of the timing noise and a comprehensive understanding of how different timing noise models ought to be included in the HMM, a challenge which is faced by other glitch detection schemes (Chukwude and Urama, 2010; Espinoza et al., 2014; Singha et al., 2021a). Independent of these timing noise considerations, future observing campaigns can provide stricter upper limits with higher cadence and/or more sensitive observations. We show that glitch recovery has a mild effect on the upper limits; the upper limits for 87% of the objects increase by no more than a factor of 5, if it is assumed that the undetected glitch recovers completely on a fiducial time-scale of 100 d.

Understanding the completeness of glitch catalogues is essential to falsifying models of glitching behaviour. For instance, Melatos et al. (2018) predicted (under certain weak, astrophysics-independent assumptions) that pulsars with large values of $-\dot{f}\langle\Delta t\rangle$ (where $\langle\Delta t\rangle$ is the mean waiting time between glitches) should show significant correlations between the size of a glitch and the forward waiting time to the next glitch. Similarly, pulsars with small values of $-\dot{f}\langle\Delta t\rangle$ are predicted to be the most likely to exhibit correlations between glitch size and backward waiting time, although the latter correlations are predicted to be weaker. As a foretaste of what is possible, Table 3.7 shows the five highest and five lowest values of $-\dot{f}\langle\Delta t\rangle$ amongst the pulsars which have been observed in the first UTMOST data release, as well as the Pearson correlation coefficients between $\Delta f/f$ and the forward and backward waiting times. While most of these pulsars do not yet have enough glitches observed for any statistically significant conclusions to be drawn, there is tentative evidence for strong forward correlations in the pulsars with the largest values of $-\dot{f}\langle\Delta t\rangle$. Continued high-cadence monitoring of these pulsars and a good understanding of the likelihood that a glitch of a given size might not have been detected are essential to falsifying the proposed relations. Improved understanding of the completeness of glitch catalogues is also important to studies of the physical conditions

involved. Several authors have investigated the nature of the neutron superfluid in the inner crust by studying glitching behaviour (e.g. Andersson et al. (2012), Ho et al. (2015), Montoli et al. (2020), and Montoli et al. (2021)). The quantity of interest is frequently the cumulative fractional change in spin frequency due to glitches, $\mathcal{A} \propto \sum_i \Delta f_i / f$. An understanding of the completeness of the glitch sample is essential to understanding the uncertainty on \mathcal{A} , and by extension understanding the implications of measured values of \mathcal{A} on the underlying physics.

Acknowledgements

The authors are most grateful to the UTMOST team for their hard work in collecting the data which forms the first UTMOST open data release. The Molonglo Observatory is owned and operated by the University of Sydney with support from the School of Physics and the University. Parts of this research are supported by the Australian Research Council (ARC) Centre of Excellence for Gravitational Wave Discovery (OzGrav) (project number CE170100004) and ARC Discovery Project DP170103625. L. Dunn is supported by an Australian Government Research Training Program Scholarship and by the Rowden White Scholarship. This work was performed on the OzSTAR national facility at Swinburne University of Technology. The OzSTAR program receives funding in part from the Astronomy National Collaborative Research Infrastructure Strategy (NCRIS) allocation provided by the Australian Government.

Data availability

The public UTMOST data underlying this work are available at <https://github.com/Molonglo/TimingDataRelease1/>. The other data underlying this work will be shared on reasonable request to the corresponding author.

3.A1 | Follow-up analysis of vetoed candidates

In Section 3.4 we describe the vetoing of three candidates. Given that these vetos proceed by removing the ToAs bracketing each of the candidates, there is a chance that genuine glitch events may be discarded by this procedure, if the glitch is close to the limit of detectability. In this appendix we investigate each of the vetoed candidates in more detail, with the aim of clarifying their origin.

3.A1.1 | PSR J0742–2822

The case of PSR J0742–2822 turns out to be straightforward. The dispersion measure (DM) in the ephemeris file provided in the UTMOST data release was mistakenly quoted as 681 pc cm^{-3} , whereas the correct value is 74 pc cm^{-3} . A large error in the DM can produce significant scatter in the timing residuals, as even small variations in the central frequency of each observation lead to large corrections to the ToAs. After correcting the DM in the ephemeris, we find that the scatter in the residuals in the vicinity of the

candidate is reduced by a factor of 20. Finally, we re-run the HM analysis with the updated DM value and find no candidate. Hence we reject the initial candidate as non-astrophysical.

3.A1.2 | PSR J1105–6107

Inspection of the timing residuals surrounding the candidate in PSR J1105–6107 reveals no obvious features, and inspection of the raw archives similarly reveals no disturbance in the vicinity of the candidate. As a consistency check, we re-generate the ToAs for PSR J1105–6107 using `PSRCHIVE` (Hotan et al., 2004) and re-run the HMM analysis. No candidate is returned, and we thus reject the candidate as non-astrophysical.

We are unable to identify a clear reason for the discrepancy between the ToAs in the data release and the re-generated set. As we are aware, both the archives containing the folded observations and the standard profile used to generate the times of arrival have not changed between the UTMOST data release and the re-generation described here. However, records of the UTMOST data release preparation are not sufficiently detailed to allow us to check this, and it is possible that the profile used to generate the ToAs in the UTMOST data release was not optimal (e.g. not sufficiently smoothed).

3.A1.3 | PSR J1359–6038

Inspection of the timing residuals surrounding the candidate in PSR J1359–6038 reveals that a single ToA at MJD 58190.7 is displaced away from the rest of the surrounding ToAs by approximately 0.4 ms. This is significant compared to the uncertainty on this ToA of $60 \mu\text{s}$. To check whether this displaced ToA is due to conditions at the observatory, we inspect the timing residuals of other pulsars that were observed no more than twelve hours before or after the ToA in question. We identify multiple pulsars in which the observation nearest MJD 58190.7 is displaced by approximately the same amount in the same direction, for example PSRs J1146–6030, J1600–3053⁹, and J1644–4559. Thus we conclude that the candidate in PSR J1359–6038 is due to local conditions at the observatory and has no astrophysical origin.

3.A2 | Synthetic dataset generation and upper limit estimation

In order to set frequentist upper limits we perform injection studies for each pulsar. We first outline the procedure for generating a single synthetic dataset for one object, which is based on a given UTMOST dataset (i.e. with identical ephemeris, observing cadence, and ToA uncertainties), with a glitch of size Δf injected.

1. A glitch epoch t_g is chosen at random, uniformly distributed between the second and second-last ToAs. Glitches which occur in either the first or last ToA gap are indistinguishable from a single outlier ToA due to some external factor, so we do not consider them when setting upper limits here.

⁹In the case of PSR J1600–3053 the displaced ToA was removed manually during the preparation of the public data release.

2. A new phase model is generated by LIBSTEMPO which matches the UTMOST phase model, except that a glitch term $\Delta\phi_g(t) = \Theta(t - t_g)\Delta f(t - t_g)$ is added [where $\Theta(t)$ is the Heaviside step function].
3. Using the ToAs of the original dataset as a starting point, a new set of ToAs is generated by shifting the original ToAs slightly so that they show zero residuals with respect to the new phase model.
4. Noise is introduced into the new set of ToAs at the levels reported in the UTMOST data release. We use the `add_efac`, `add_equad`, and `add_rednoise` functions in LIBSTEMPO, using the EFAC, EQUAD, and red noise parameters reported for each pulsar in the UTMOST data release.
5. The new phase model and new set of ToAs are saved as a synthetic dataset.

This procedure ensures that the synthetic datasets closely match the true datasets in various important aspects, e.g. basic timing model parameters, observing cadence, and noise characteristics.

The procedure for estimating the probability of detection for a glitch size Δf [denoted $P_d(\Delta f)$] in a single pulsar is straightforward:

1. Generate 100 synthetic datasets with a glitch of size Δf injected, according to the procedure in the paragraph above.
2. Analyse each dataset with the HMM following the method outlined in Section 3.3.
3. For each dataset, determine the Bayes factor K between the model $M_1(k_{\text{inj}})$ with a glitch included in the ToA gap indexed by k_{inj} corresponding to the injected glitch epoch and the model M_0 with no glitch included.
4. The proportion of synthetic datasets with $K > K_{\text{th}}$ gives an estimate of $P_d(\Delta f)$.

Finally we give a simple prescription for estimating the value of $\Delta f^{90\%}$ from $P_d(\Delta f^{90\%}) = 0.9$, i.e. the 90% upper limit on the size of undetected glitches in each pulsar.

1. Choose the starting range of glitch sizes to be $[\Delta f_-, \Delta f_+] = [10^{-9}, 10^{-6}]$ Hz.
2. Choose a glitch size Δf by bisecting the range logarithmically, i.e.,

$$\log_{10}(\Delta f) = \log_{10}(\Delta f_-) + [\log_{10}(\Delta f_+) - \log_{10}(\Delta f_-)]/2 \quad (3.13)$$

where all the frequencies are understood to be in units of Hz.

3. Calculate $P_d(\Delta f)$ as outlined previously.
4. If $|P_d(\Delta f) - 0.9| \leq 0.01$, terminate and take Δf as the 90% frequentist upper limit $\Delta f^{90\%}$.
5. Otherwise, revise the glitch size range as follows:
 - (a) If $P_d(\Delta f) > 0.9$, set $\Delta f_+ = \Delta f$.
 - (b) If $P_d(\Delta f) < 0.9$, set $\Delta f_- = \Delta f$.
6. Return to step (ii).

This is essentially a binary search over possible upper limits.

3.A3 | Posterior distributions and frequency tracks for HMM analyses

This appendix collects Figs. 3.4–3.11 showing the sequence of most likely frequencies $\hat{f}(t_n)$ and heatmaps of the frequency posterior distributions $\gamma_f(t_n)$ for the eight pulsars which are followed up with glitch parameter estimation analyses as described in Section 3.4. The structure of each figure is essentially the same: the left panel shows $\hat{f}(t_n)$ as a function of the MJD, and the right panel shows $\ln[\gamma_f(t_n)]$ as a function of ToA gap index. The values of $\ln[\gamma_f(t_n)]$ have been clipped below to aid readability. The vertical axis in both cases extends over the full f range in the DOI for each analysis. In some cases $\gamma_f(t_n)$ displays multiple peaks; see Section 3.4.1 and (Dunn et al., 2021a) for further discussion of this phenomenon. Both $\hat{f}(t_n)$ and $\gamma_f(t_n)$ are obtained using the forward-backward algorithm (Rabiner, 1989). We also remind the reader that $\hat{f}(t_n)$ is the sequence of most likely states at each timestep (i.e. it is constructed from the sequence of modes of the posterior distribution of states) – it is not the most likely sequence of states, which may instead be calculated using the Viterbi algorithm (Rabiner, 1989). However, the difference between these two sequences is typically small (Melatos et al., 2020), and so we prefer to use $\hat{f}(t_n)$.

We also include in Fig. 3.12 an illustrative plot showing the frequency derivative posterior distribution $\gamma_{\dot{f}}(t_n)$ from the follow-up analysis of PSR J1371–3744. As mentioned in Section 3.4, we do not include equivalent plots for every pulsar, as the coarse discretisation of \dot{f} in the DOI leads to relatively uninformative \dot{f} posteriors. In the exemplar plot, $\gamma_{\dot{f}}(t_n)$ shows support over a significant fraction of the DOI, particularly after the glitch occurs at the 103rd ToA, making it difficult to make useful inferences about the evolution of \dot{f} over the dataset. Note that the heatmap shows $\gamma_{\dot{f}}$, not its natural logarithm, unlike Figs. 3.4–3.11.

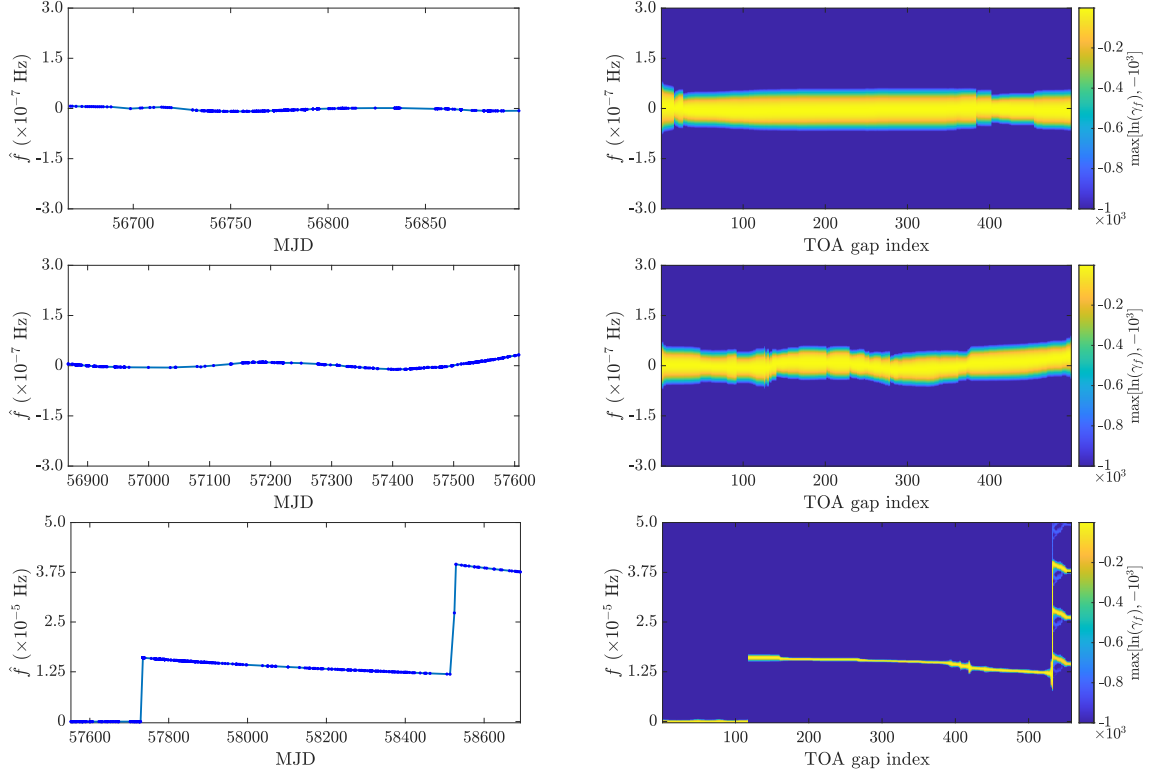


Figure 3.4: Sequence of most likely frequencies $\hat{f}(t_n)$ (*left*) and heatmap of posterior frequency probability $\ln[\gamma_f(t_n)]$ (*right*) for the HMM follow-up analysis of PSR J0835–4510. Frequency is on the vertical axis in all panels, and the range of the vertical axis is the full range of the DOI. Note that the horizontal axes for the two panels are not exactly the same: the left panels have MJD on the horizontal axis, while the right panels have ToA gap index on the horizontal axis instead, for ease of plotting. The three rows correspond to sections 1, 2 and 3 from top to bottom as described in Table 3.5.




Figure 3.5: As in Figure 3.4, but for PSR J1257–1027.




Figure 3.6: As in Figure 3.4, but for PSR J1452–6036.




Figure 3.7: As in Figure 3.4, but for PSR J1622–4950.




Figure 3.8: As in Figure 3.4, but for PSR J1703–4851.




Figure 3.9: As in Figure 3.4, but for PSR J1709–4429.




Figure 3.10: As in Figure 3.4, but for PSR J1731–4744.




Figure 3.11: As in Figure 3.4, but for PSR J1740–3015.




Figure 3.12: Heatmap of posterior frequency derivative probability $\gamma_f(t_n)$ for the HMM follow-up analysis of J1731–4744.

4

A new small glitch in Vela discovered with a hidden Markov model

This chapter is a reproduction of Dunn et al. (2023a), *Monthly Notices of the Royal Astronomical Society* 522.4, pp. 5469–5478.

Abstract

A striking feature of the Vela pulsar (PSR J0835–4510) is that it undergoes sudden increases in its spin frequency, known as glitches, with a fractional amplitude on the order of 10^{-6} approximately every 900 days. Glitches of smaller magnitudes are also known to occur in Vela. Their distribution in both time and amplitude is less well constrained but equally important for understanding the physical process underpinning these events. In order to better understand these small glitches in Vela, an analysis of high-cadence observations from the Mount Pleasant Observatory is presented. A hidden Markov model (HMM) is used to search for small, previously undetected glitches across 24 years of observations covering MJD 44929 to MJD 53647. One previously unknown glitch is detected around MJD 48636 (Jan 15 1992), with fractional frequency jump $\Delta f/f = (8.19 \pm 0.04) \times 10^{-10}$ and frequency derivative jump $\Delta \dot{f}/\dot{f} = (2.98 \pm 0.01) \times 10^{-4}$. Two previously reported small glitches are also confidently re-detected, and independent estimates of their parameters are reported. Excluding these events, 90% confidence frequentist upper limits on the sizes of missed glitches are also set, with a median upper limit of $\Delta f^{90\%}/f = 1.35 \times 10^{-9}$. Upper limits of this kind are enabled by the semi-automated and computationally efficient nature of the HMM, and are crucial to informing studies which are sensitive to the lower end of the glitch size distribution.

4.1 | Introduction

The electromagnetic spindown of pulsars is sometimes interrupted by a glitch – a sudden increase in the spin frequency. At the time of writing, 670 glitches among 208 pulsars have been recorded in the Jodrell Bank Observatory (JBO) catalogue¹ (Espinoza et al., 2011). Glitches provide a valuable window into the interior physics of neutron stars, but the underlying physical cause of glitches remains unknown; see Haskell and Melatos (2015) for a review of proposed mechanisms.

One of the most prolific glitching pulsars is the Vela pulsar (PSR J0835–4510), which has 24 recorded glitches since its discovery in 1968 (Large et al., 1968). Historically it has been a target of extensive monitoring (e.g. Cordes et al. (1988) and Dodson et al. (2007)) and many studies of glitching behaviour [e.g. Dodson et al. (2002), Melatos et al. (2008), Shannon et al. (2016), Palfreyman et al. (2018), Fuentes et al. (2019), and Espinoza et al. (2021)]. Vela has been monitored extensively by the Mount Pleasant Observatory near Hobart, Australia, accumulating roughly two decades of high-cadence timing data between 1981 and 2005 (Dodson et al., 2007). This has proved to be a rich source of information on the timing behaviour of Vela, revealing several large glitches, including high time resolution observations of the large glitch of January 2000 (Dodson et al., 2002). The high density of the observations also enables the identification of more subtle timing features. Using later single-pulse observations of the 2016 Vela glitch from Mount Pleasant (Palfreyman et al., 2018), Ashton et al. (2019a) related the observed, short-term, post-glitch recovery to the dynamics of three coupled components in the stellar interior. Separately, Espinoza et al. (2021) presented a systematic search for small glitches and other irregularities in Vela’s rotation, finding two new small glitches as well as a population of small-amplitude changes in the spin frequency of both signs, which are regarded as a component of the timing noise (Cordes et al., 1988; D’Alessandro et al., 1995; Chukwude and Urama, 2010).

While timing data from Vela have been investigated extensively over the years, it is not common practice to employ Monte Carlo simulations to quantify the completeness of the glitch catalogues reported based on these data (Janssen and Stappers, 2006; Yu and Liu, 2017). In this paper, we use a hidden Markov Model (HMM) to address this gap (Melatos et al., 2020)². The method has previously been successfully used to correct glitch parameter values whose estimation was confounded by periodic scheduling (Dunn et al., 2021a), and to set upper limits on the size of undetected glitches in datasets from timing programmes undertaken at both Parkes (Lower et al., 2021) and UTMOST (Dunn et al., 2022b).

We provide an overview of the HMM glitch detector in Section 4.2. Upon conducting an automated HMM search for missed glitches between Nov 21 1981 and Oct 24 2005, we discover one small new glitch, discussed in Section 4.3. In Section 4.4 we present re-detections and re-analyses of the two small glitches previously detected by Espinoza et al. (2021). Finally, we put systematic frequentist upper limits on the size of any further missed glitches in our Vela dataset, as discussed in Section 4.5.

¹<http://www.jb.man.ac.uk/pulsar/glitches/gTable.html>

²https://github.com/ldunn/glitch_hmm

4.2 | HMM glitch detector

In this section we give a brief outline of the HMM glitch detection algorithm. The reader is referred to Melatos et al. (2020) for a complete description, and to Rabiner (1989) for an overview of HMMs.

4.2.1 | Pulsar timing in an HMM framework

HMMs assume that the system in question has a well-defined internal state at each timestep, which evolves according to some Markov process. However, the internal state is “hidden”; we rely on some proxy measurement which implies a probability of occupying each hidden state. In this paper the system is the Vela pulsar, and the internal state is the combination of its spin frequency f and frequency derivative \dot{f} (discretised on a grid). The measurements are the barycentred pulse times of arrival (ToAs), which are connected to the underlying hidden states by the requirement that the number of rotations accumulated between two consecutive ToAs should be an integer (allowing for small deviations away from this requirement due to uncertainty on the ToAs and the discretisation of f and \dot{f}). This requirement is represented in the HMM through the “emission probability”, which expresses the probability of observing a particular measurement (namely the gap between consecutive ToAs, denoted by z) given the system is in a particular hidden state (f and \dot{f} pair, denoted by q_i). Mathematically, we write the emission probability $\Pr(z | q_i)$ in terms of a von Mises distribution

$$\Pr(z | q_i) = \frac{\exp[\kappa \cos(2\pi\Phi)]}{2\pi I_0(\kappa)}, \quad (4.1)$$

where Φ denotes the number of turns accumulated during the ToA gap z given an (f, \dot{f}) pair, κ is a parameter containing the uncertainty in the phase due to both ToA uncertainty and the discretisation of f and \dot{f} , and $I_0(x)$ is the zeroth modified Bessel function of the first kind. Explicit expressions for Φ and κ as a function of q_i are given in equations (1) and (2) of Dunn et al. (2022b), and the mathematical form of $\Pr(z | q_i)$ given in equation (4.1) is justified in Melatos et al. (2020).

The Markov process which drives the evolution of the hidden state is assumed in this work to be a random walk in the frequency derivative, i.e.

$$\frac{d^2 f}{dt^2} = \xi(t), \quad (4.2)$$

where $\xi(t)$ is a white noise term satisfying

$$\langle \xi(t) \rangle = 0, \quad (4.3)$$

$$\langle \xi(t)\xi(t') \rangle = \sigma_{\text{TN}}^2 \delta(t - t'), \quad (4.4)$$

where σ_{TN} is a free parameter which determines the strength of the timing noise included in the model. Equation (4.2) applies in the interval between glitches. At any instant where a glitch occurs, f and \dot{f} jump discontinuously by Δf_p and $\Delta \dot{f}_p$ respectively (where p denotes permanent; see Section 4.2.3). The no-glitch ephemeris generated by

traditional timing methods provides an approximation to the initial conditions satisfied by equation (4.2).

In the above approach to pulsar timing we track the underlying spin evolution based on measurements of the phase locally, as inferred from consecutive ToAs, rather than construct a model of the phase globally. Approaching this task using an HMM allows us to take advantage of existing algorithms to efficiently compute useful quantities such as the posterior distributions on f and \dot{f} , as well as the probability of an observed data set, given a particular model (i.e. the Bayesian evidence). The computational framework provided by the HMM formalism is a central motivation for adopting this approach.

4.2.2 | Glitch detection with the HMM

As mentioned in the previous section, the structure of the HMM makes it computationally efficient to compute the model evidence $\Pr(D | M)$ where D is the data analysed and M is some particular setup of the HMM. The algorithm used to perform this calculation is known as the forward algorithm (Rabiner, 1989; Melatos et al., 2020). In particular, we can specify two classes of models. In one, denoted M_0 , the evolution of the pulsar's rotation is governed solely by a combination of secular spindown and timing noise. In the other, denoted $M_n(k_1, k_2, \dots, k_n)$, the evolution of the pulsar's rotation is governed by the secular spindown and timing noise *except* during n ToA gaps containing glitches indexed by $\{k_1, k_2, \dots, k_n\}$. The characteristics of the glitch are not specified by the model except that the frequency jump Δf_p is assumed to be instantaneous and constrained to be positive. An instantaneous jump in frequency derivative $\Delta \dot{f}_p$ is also allowed, and may be of either sign.

With the above setup, model selection proceeds by computing the Bayes factors

$$K_1(k) = \frac{\Pr[D | M_1(k)]}{\Pr[D | M_0]}. \quad (4.5)$$

These $K_1(k)$ are compared against a detection threshold K_{th} which is chosen ahead of time according to the analyst's preference. Here we follow Melatos et al. (2020) and Dunn et al. (2022b) in choosing $\ln(K_{\text{th}}) = 1.15$ (corresponding to $K_{\text{th}} = 10^{1/2}$). The possibility of multiple glitches in a dataset is handled via the greedy procedure described in section 4.2 of Melatos et al. (2020). Briefly, if at least one $K_1(k)$ exceeds K_{th} , we choose

$$k^* = \arg \max_{2 \leq k \leq N_T - 1} K_1(k) \quad (4.6)$$

and proceed to calculate the Bayes factors

$$K_2(k^*, k_2) = \frac{\Pr[D | M_2(k^*, k_2)]}{\Pr[D | M_1(k^*)]} \quad (4.7)$$

for $2 \leq k_2 \neq k^* \leq N_T - 1$. This procedure repeats until none of the computed Bayes factors exceed K_{th} .

4.2.3 | Glitch measurement with the HMM

In order to measure the glitch properties using the HMM, we require estimates of the value of the hidden state as a function of time. To do this, we use the forward-backward




Figure 4.1: Illustrative heatmap showing the hidden state posterior $\gamma_{q_i}(t_n)$ for a single timestep in the analysis of the glitch event described in Section 4.3. The color scale reflects the posterior density at the 85th timestep, $\gamma_{q_i}(t_{85})$. The plotted range in f and \dot{f} reflects the complete allowed space of hidden (f, \dot{f}) states in this analysis, and the hidden state posterior is well-localised within this range.

algorithm (Rabiner, 1989; Melatos et al., 2020), an extension of the forward algorithm which is used to calculate the model evidences, to obtain the posterior distribution on the hidden state $q_i = (f, \dot{f})$ at each timestep. This distribution is denoted by $\gamma_{q_i}(t_n)$. The t_n are the timestamps associated with the observed data used in this analysis: the barycentred times of arrival. Thus we emphasise that $\gamma_{q_i}(t_n)$ is evaluated at a discrete set of times typically separated by one to a few days. An exemplar heatmap of $\gamma_{q_i}(t_n)$ is shown in Figure 4.1. From this posterior we construct the one-dimensional posteriors for f and \dot{f} by marginalising over \dot{f} and f , respectively denoted $\gamma_f(t_n)$ and $\gamma_{\dot{f}}(t_n)$.

We take the mode of $\gamma_f(t_n)$ at each timestep to construct the sequence of most likely frequencies during each ToA gap, denoted $f^*(t_n)$. An analogous procedure is used to construct the sequence of most likely frequency derivatives, $\dot{f}^*(t_n)$. For the purposes of parameter estimation, our model for the phase induced by a glitch includes a permanent (indicated by a subscript p) step change in phase³, frequency, and frequency derivative at the epoch of the glitch t_g , as well as an exponentially decaying component (indicated by the subscript d):

$$\Delta\phi_g(t) = \Theta(t - t_g) \left\{ \Delta\phi + \Delta f_p(t - t_g) + \frac{1}{2} \Delta \dot{f}_p(t - t_g)^2 + \tau_d \Delta f_d \left[1 - e^{-(t-t_g)/\tau_d} \right] \right\}. \quad (4.8)$$

³The step change in phase in equation (4.8) does not represent a physical process. Rather it allows us to account for an incorrect estimate of the glitch epoch t_g in a manner which is relatively easy to fit. A physical step change in phase may occur if, for example, there is a change in the magnetic longitude of the pulsar’s radio beam due to magnetospheric processes at the time of the glitch, e.g. Palfreyman et al. (2018), but we do not consider such processes here.

In equation (4.8), $\Theta(x)$ denotes the Heaviside step function. With the exception of the glitch described in Section 4.4.2, we set $\Delta f_d = 0$. To estimate the glitch parameters from the HMM output we fit the time derivative of equation (4.8) to the post-glitch values of $f^*(t_n)$ using the LMFIT package (Newville et al., 2014). The $f^*(t_n)$ values are assigned uncertainties based on the width of $\gamma_f(t_n)$, estimated as the standard deviation of a Gaussian distribution fitted to $\gamma_f(t_n)$. These Gaussian fits are also performed using LMFIT.

The HMM differs from the traditional method of fitting a global phase model to the ToAs. Although the HMM incorporates phase information from the ToAs directly through $\Pr(z | q_i)$ (see Section 4.2.1), it returns information about the evolution of f and \dot{f} , rather than parameters of a phase model which applies over the entire dataset. Hence if we are to estimate the parameters of the glitch using the HMM we are obliged to model the post-glitch frequency evolution rather than phase evolution. Glitch parameter estimation based on the values of $f^*(t_n)$ and $\dot{f}^*(t_n)$ has been performed in previous analyses (Dunn et al., 2021a; Dunn et al., 2022b). The results therein as well as in Sections 4.3 and 4.4 of this work return glitch parameters broadly in line with those returned by traditional TEMPO2 or TEMPONEST-based analyses. In all cases the parameter estimates should be understood as conditional on the model assumed. In the case of the HMM this includes both the HMM’s model of the dynamics of f and \dot{f} [e.g. equation (4.2)] as well as the model of post-glitch frequency evolution used [viz. equation (4.8)].

4.3 | A small glitch at MJD 48636

Here we describe the detection and subsequent analysis of a new small glitch in the Vela pulsar. The data analysed are described in sections 2 and 3 of Espinoza et al. (2021). In particular we analyse the daily effective times of arrival described in section 3.2, rather than the full dataset which contains hundreds to thousands of ToAs per day. Using the full dataset would inflate the computational cost but would not be expected to enhance the sensitivity of the HMM glitch detector; see Section 4.5.2 for a brief discussion of the effect of observing cadence on sensitivity. In total the data cover MJD 44929 to MJD 53667, and are broken down into smaller sections, typically using either known glitches or large observing gaps as boundaries. Where possible, the data sections are allowed to overlap somewhat to mitigate the chance of missing a glitch which lies close to a section boundary.

For each section we use TEMPO2 to fit values of f , \dot{f} and \ddot{f} to incorporate as an initial condition the local, secular evolution of the pulsar in the HMM analysis. In a few cases where the data section begins immediately after a large glitch and its associated recovery, we exclude the first ~ 50 days of data to avoid the majority of the recovery, which the HMM is unable to handle without artificially inflating the strength of the timing noise tracked by the HMM [see equation (4.12)]. Shannon et al. (2016) found evidence for a common recovery timescale in PSR J0835–4510 of 25 days (in addition to a much longer timescale of 1300 days which the HMM is not confounded by), so by excluding 50 days we expect to exclude approximately 86% of the recovery (i.e. two e -foldings). Details of the data sectioning are given in Table 4.1.

| Section | MJD start | MJD end |
|---------|-----------|---------|
| 1 | 44964 | 45191 |
| 2 | 45211 | 45446 |
| 3 | 45421 | 46257 |
| 4 | 46285 | 46562 |
| 5 | 47894 | 48049 |
| 6 | 48030 | 48456 |
| 7 | 48475 | 48549 |
| 8* | 48551 | 48719 |
| 9 | 48680 | 49558 |
| 10 | 49560 | 49590 |
| 11 | 49594 | 49919 |
| 12 | 49882 | 50367 |
| 13 | 50370 | 50703 |
| 14 | 51294 | 51424 |
| 15 | 51425 | 51558 |
| 16 | 51559 | 52039 |
| 17 | 52001 | 52339 |
| 18 | 53095 | 53192 |
| 19 | 53194 | 53218 |
| 20 | 53223 | 53500 |
| 21 | 53528 | 53667 |

Table 4.1: Details of the data sectioning used for the HMM glitch search. Horizontal rules indicate the location of known glitches. The asterisk indicates the data section containing the newly discovered glitch described in Section 4.3.

4.3.1 | Detection of the new glitch

As discussed in Section 4.2.2, the HMM detects glitches by Bayesian model selection, comparing models which contain a glitch in a given ToA gap to the model which contains no glitches.

In analysing the stretch of data between MJD 48551 and MJD 48719, the HMM returns a glitch candidate between the ToAs at MJD 48635.6 and MJD 48636.6, with a log Bayes factor $\ln[K_1(k^*)] = 2.4$ ($k^* = 81$). Fig. 4.2 shows $\ln[K_1(k)]$ for each ToA gap, indexed by k . A clear peak is visible at $k^* = 81$. The HMM is sensitive to even single ToAs which are significantly displaced from the trend. Therefore we check that the glitch candidate is not due to a non-astrophysical disturbance of the surrounding ToAs by re-running the analysis with the two ToAs bracketing the glitch candidate removed. A log Bayes factor of 2.9 is returned, confirming that the candidate is not simply due to a disturbed ToA in the immediate vicinity of the glitch. The veto procedure of removing the bracketing ToAs is necessary because the HMM cannot itself distinguish between a ToA which is in error due to radio frequency interference (RFI), or an observation with low signal-to-noise ratio, and so on, and a ToA which is displaced because of a glitch. The HMM treats all ToAs equally. In accommodating a displaced ToA it may be that the HMM inference prefers a glitch-containing model which “explains” the displaced ToA via some change in (f, \dot{f}) over the model with no glitch, and hence a spurious glitch candidate is returned. Furthermore, due to the inclusion of spin-wandering in the HMM via equation (4.2), an incorrect model which includes a spurious change in (f, \dot{f}) is free to wander back to the true spin evolution over the longer term. Hence the inclusion of data taken long after the RFI event does not significantly penalise the incorrect model. This is in contrast to a model which does not include spin wandering, where a change in (f, \dot{f}) induces phase residuals which grow monotonically over time. We expect t_g to be either immediately before or after a displaced ToA, so re-running the search with the ToAs bracketing the glitch candidate removed eliminates the RFI scenario.

We also check that dispersion measure (DM) effects are not masquerading as a glitch. We measure the evolution of the DM over the data span containing the glitch candidate using the high-cadence ToAs at 635 and 950 MHz [with a cadence of approximately 2 minutes at both observing frequencies covering about 70% of each day (Espinoza et al., 2021)], looking for sharp changes in the DM around the epoch of the glitch candidate. No such changes are apparent and we find that a linear model with a DM of $68.288 \text{ pc cm}^{-3}$ at MJD 48616 and a slope of $-0.036 \text{ pc cm}^{-3} \text{ yr}^{-1}$ describes the data well. This is a minor change from the original DM model which is constant at 68.29 pc cm^{-3} for the section of data containing the glitch candidate. Re-analysing the data with the refined, linearly-sloping DM model increases the log Bayes factor of the candidate marginally to 2.5, and we use the refined DM model in the rest of the analysis. Finally, we check for further glitches in this stretch of data following the greedy procedure outlined in Section 4.2.2. None of the $K_2(k^*, k_2)$ exceed K_{th} .

The next most significant glitch candidate over the full 24 years of data has $\ln[K_1(k)] = 0.3$. The mean value of $\ln[K_1(k)]$ across the full 24 years (excluding the stretch of data containing MJD 48636) is $\langle \ln[K_1(k)] \rangle = -2.21$, with a standard deviation of 1.05. The glitch candidate at MJD 48636 is thus a clear outlier, and unlikely to simply be a timing noise excursion. The event at MJD 48636 was also flagged in the analysis performed by Espinoza et al. (2021), but was ultimately discarded. Its characteristics, and the grounds




Figure 4.2: Sequence of $\ln [K_1(k)]$ for each ToA gap in the section of data containing the glitch described in Section 4.3. A clear peak is visible at the 81st ToA gap, corresponding to MJD 48635.6 – MJD 48636.6. The threshold $\ln(K_{\text{th}}) = 1.15$ is shown as the dashed line.

on which it was discarded, are summarised in Appendix 4.A1.

4.3.2 | Parameter estimation

With the detection of the new glitch established, we move on to estimate its parameters, following the procedure outlined in Section 4.2.3. The key outputs from the HMM for this purpose are the marginalised posteriors of f and \dot{f} , heatmaps of which are shown in the top two panels of Fig. 4.3. We note two features of these heatmaps, which also apply to the analogous figures for the glitch events described in Section 4.4, namely Figs. 4.5 and 4.6. First, the value shown on the heatmaps between t_n and t_{n+1} is the value of the posterior distribution evaluated at t_n . For long ToA gaps this can give the impression that the frequency remains constant for many days in certain parts of the post-glitch heatmap, for example between MJD 48650 and MJD 48660 in the top left panel of Fig. 4.3. We prefer for transparency to have this aspect of $\gamma_{q_i}(t_n)$ reflected in the heatmaps, rather than choose an interpolation scheme which artificially evolves $\gamma_{q_i}(t_n)$ over a single ToA gap to give the impression of smooth f and \dot{f} evolution. Second, the posteriors visibly widen at either end of the data span, and close to the epoch of the glitch t_g . This is a real effect, because at those times the evolution of the hidden state is less constrained. Typically the posterior at a given time t_n is constrained by the data preceding and following t_n , and data which are further separated from t_n are less constraining. However, the posterior at $t_n < t_g$ is not constrained by the data at $t_n > t_g$, because of the considerable freedom in both f and \dot{f} evolution allowed at t_g . The same argument applies to posteriors calculated at $t_n > t_g$, which are not constrained by the data at $t_n < t_g$.

We also note the small fluctuations in $\dot{f}^*(t_n)$ which are visible in the bottom right panel of Fig. 4.3. These are real in the sense that they are not visual artefacts, and correspond

| Parameter (units) | HMM fit (via LMFIT) | TEMPONEST fit |
|------------------------------------|------------------------------------|--------------------------------------|
| t_g (MJD) | 48635.64 | 48635.64 |
| $\Delta\phi$ (turns) | N/A | $(-1.5 \pm 0.8) \times 10^{-3}$ |
| Δf_p (Hz) | $(9.24 \pm 0.03) \times 10^{-9}$ | $(9.3 \pm 0.8) \times 10^{-9}$ |
| $\Delta \dot{f}_p$ (Hz s $^{-1}$) | $(-4.67 \pm 0.01) \times 10^{-15}$ | $-3.2_{-1.6}^{+1.7} \times 10^{-15}$ |

Table 4.2: Glitch model parameters returned by the HMM (via LMFIT) and TEMPONEST for the new small glitch described in Section 4.3.2. In both cases the glitch epoch is taken to be immediately after the last pre-glitch ToA, to match the HMM. With TEMPONEST we do not fit for t_g , instead fitting for the unphysical parameter $\Delta\phi$. The given ranges are the 90% confidence intervals in the LMFIT case, and 90% credible intervals in the TEMPONEST case.

to genuine changes in the mode of $\gamma_j(t_n)$ by one \dot{f} bin at a time. However, given that the scale of the fluctuations is much smaller than the width of $\gamma_j(t_n)$, we do not ascribe any physical significance to these features.

Parameter estimation with the HMM proceeds by fitting the maximum *a posteriori* values of the spin frequency, $f^*(t_n)$, against the model for the post-glitch frequency evolution given by the time derivative of equation (4.8). Note that when performing the latter analysis, the values of f , \dot{f} and \ddot{f} are fitted using only the pre-glitch data. This is in contrast to the HMM analyses searching for undetected glitches where the values are derived using the complete data in each section (at this stage, there is no notion of pre- and post-glitch, as a glitch has not yet been detected). Fitting to the pre-glitch data ensures that the derived $f^*(t_n)$ and $\dot{f}^*(t_n)$ tracks are approximately flat before the glitch occurs, which simplifies the parameter estimation.

As a control experiment, we also estimate the glitch parameters using TEMPONEST (Lentati et al., 2014), which includes both deterministic terms [f , \dot{f} , \ddot{f} , and the glitch parameters $\Delta\phi$, Δf_p , and $\Delta \dot{f}_p$ in equation (4.8)] and a model of the phase residuals due to timing noise as a red noise process with power spectral density

$$P(f) = \frac{A^2}{12\pi^2} \left(\frac{f}{f_{\text{yr}}} \right)^{-\beta}, \quad (4.9)$$

where $f_{\text{yr}} = (1 \text{ yr})^{-1}$.

The results of the HMM and TEMPONEST fits are reported in Table 4.2. The TEMPO2 phase residuals before and after fitting for the glitch with TEMPONEST are shown in Fig. 4.4. Note that the post-fit residuals in Fig. 4.4 are not flat. The TEMPONEST phase model includes timing noise as a sum of sinusoids (Lentati et al., 2013), and the sinusoids are not subtracted from the phase when forming residuals in TEMPO2. Hence the sinusoidal structure in the post-fit residuals reflects the analytic form of the timing noise assumed in the TEMPONEST analysis. The rms of the post-fit residuals is 2.5×10^{-4} s. To compare this value with the value expected from the TEMPONEST results, we assume that the residuals from TEMPONEST are dominated by the lowest-frequency component of the red noise process, which has frequency T^{-1} where $T = 0.459 \text{ yr}$ is the total timespan of the data section containing the glitch. Given the returned red noise parameters $A = 10^{-8.8} \text{ yr}^{3/2}$,




Figure 4.3: (*Top*) Heatmaps of the marginalised posterior distributions of f (*left*) and \dot{f} (*right*) for the data section containing the newly discovered glitch described in Section 4.3.2. (*Bottom*) Sequence of *a posteriori* most probable f (*left*) and \dot{f} (*right*) states over time, for the same data section. The values of f and \dot{f} are relative to the pre-glitch timing solution, hence they are flat and centred on zero in the pre-glitch region.

$\beta = 8.3$, we expect rms residuals $\sigma \propto \sqrt{P(T^{-1})}$ on the order of

$$\sigma = \frac{1}{\sqrt{T}} \frac{A}{\sqrt{12}\pi} \left(\frac{T^{-1}}{f_{\text{yr}}} \right)^{-\beta/2}, \quad (4.10)$$

i.e. $\sigma \approx 2.7 \times 10^{-4}$ s, in agreement with the post-fit rms residuals. Equation 4.10 assumes an effective bandwidth $\sim 1/T$ to match the spacing of frequency components used in TEMPONEST.

With a fractional glitch size of $\Delta f/f = (8.19 \pm 0.04) \times 10^{-10}$ (according to the HMM estimate), this is the second smallest glitch reported in Vela.⁴

4.4 | Other small glitches

In this work we do not re-analyse formally the several large ($\Delta f/f \gtrsim 10^{-6}$) glitches which occurred during the 24-year observing span, as their existence is almost certainly beyond dispute, and is verified by a quick and informal check with the HMM. However, for the sake of completeness, and to compare with the results of previous work, we do investigate the two small glitches reported by Espinoza et al. (2021), using the HMM. This section presents the findings of the investigation. In short, both events exceed the HMM's Bayes factor threshold for detection, viz. $\ln K_{\text{th}} = 1.15$.

⁴The smallest is a glitch with $\Delta f/f = 3.9 \times 10^{-10}$ which was reported by Jankowski et al. (2015b), although subsequently Lower et al. (2020) have suggested that a timing noise-only model is favored.





Figure 4.4: Phase residuals from TEMPO2 before (*top*) and after (*bottom*) fitting for the new small glitch at MJD 48636 with TEMPNREST. The location of the glitch is marked by the vertical dotted line. Note that the post-fit residuals are not flat, because the TEMPNREST fit includes a timing noise component which is not subtracted from the residuals.

| Parameter (units) | HMM fit (via LMFIT) | TEMPONEST fit | Espinoza et al. (2021) |
|------------------------------------|----------------------------------|--------------------------------------|-----------------------------------|
| t_g (MJD) | 48549.9 | 48549.9 | 48550.37 ± 0.03 |
| $\Delta\phi$ (turns) | N/A | $(0.4 \pm 1.3) \times 10^{-3}$ | N/A |
| Δf_p (Hz) | $(5.86 \pm 0.01) \times 10^{-8}$ | $(5.7 \pm 0.1) \times 10^{-8}$ | $(6.21 \pm 0.05) \times 10^{-8}$ |
| $\Delta \dot{f}_p$ (Hz s $^{-1}$) | $(-5.6 \pm 0.1) \times 10^{-15}$ | $-0.3_{-2.2}^{+2.3} \times 10^{-15}$ | $(-16.5 \pm 0.3) \times 10^{-15}$ |

Table 4.3: Glitch model parameters returned by the HMM (via LMFIT) and TEMPONEST for the small glitch of 1991 described in Section 4.4.1. In the second and third columns the glitch epoch is taken to be immediately after the last pre-glitch ToA, to match the HMM. With TEMPONEST we fit for $\Delta\phi$ instead of t_g . The given ranges are the 90% confidence intervals in the second and fourth columns, and 90% credible intervals in the third column.

4.4.1 | A small glitch in 1991

Espinoza et al. (2021) reported a small glitch in Vela at MJD 48550.37(2) with $\Delta f = 6.21(3) \times 10^{-8}$ Hz. We first check whether the HMM re-detects this glitch, analysing data between MJD 48475 and MJD 48630. We find a single glitch candidate during the ToA gap between MJD 48549.9 and MJD 48551.9, with $\ln[K_1(k^*)] = 107.9$ ($k^* = 57$). This is a clear detection.

We then estimate the glitch parameters taking the same approach as in Section 4.3.2, using both the HMM and TEMPONEST separately. The results of these analyses are shown in Table 4.3, along with the parameter estimates from Espinoza et al. (2021). As the evolution of $f^*(t_n)$ is nonlinear over the full post-glitch timespan, we fit only the first 30 days' worth of post-glitch evolution to the HMM output, and only fit for Δf_p and $\Delta \dot{f}_p$. This approach matches the approach taken in Espinoza et al. (2021). For the TEMPONEST fits, we do not truncate the data, but nor do we fit for a $\Delta \ddot{f}_p$ term. The estimates of Δf_p and $\Delta \dot{f}_p$ from the HMM, TEMPONEST, and Espinoza et al. (2021) differ significantly. All three 90% credible (or confidence, for TEMPONEST) intervals are disjoint for both parameters. This is to be expected. The three methods take different approaches to measuring the glitch parameters and incorporating the effects of timing noise, and so the underlying models which the estimates are conditional upon differ substantially. As an example of the effect that including timing noise can have on the recovered glitch parameters, fitting with TEMPONEST but excluding timing noise from the model produces a set of glitch parameters which is nearly disjoint from the other three sets listed in Table 4.3: $\Delta f_p = (6.68 \pm 0.05) \times 10^{-8}$ Hz and $\Delta \dot{f}_p = (-2.4 \pm 0.8) \times 10^{-15}$ Hz s $^{-1}$.

4.4.2 | A small glitch in 1999

Espinoza et al. (2021) also reported the detection of a small glitch in Vela at MJD 51425.12, with a detectable exponential recovery. Analysis with the HMM of the data between MJD 51294 and MJD 51558 reveals a glitch candidate during the ToA gap between MJD 51424.8 and 51425.9 with $\ln[K_1(k^*)] = 1.4 \times 10^4$ ($k^* = 75$).

As in Section 4.4.1, we calculate the posterior distribution of the hidden state and maximum *a posteriori* tracks in f and \dot{f} , the results of which are shown in Fig. 4.5. Interestingly the HMM finds the event successfully, even though exponential post-glitch recoveries are not part of the HMM's dynamical model (Melatos et al., 2020; Dunn et al.,




Figure 4.5: (*Top*) Heatmaps of the marginalised posterior distributions of f (*left*) and \dot{f} (*right*) for the data section containing the 1991 glitch described in Section 4.4.1. (*Bottom*) Sequence of *a posteriori* most probable f (*left*) and \dot{f} (*right*) states over time, for the same data section. The values of f and \dot{f} are relative to the pre-glitch timing solution, hence they are flat and centred on zero in the pre-glitch region.

2021a). This confirms the event’s significance; an HMM model including post-glitch recoveries (a nontrivial extension which is the goal of future work) would produce an even higher Bayes factor, i.e. $\ln[K_1(k^*)] > 1.4 \times 10^4$. To accommodate the exponential decay which is visible in Fig. 4.6 we increase artificially the value of the parameter which controls the strength of the timing noise included in the HMM, σ_{TN} , to $5 \times 10^{-17} \text{ Hz s}^{-3/2}$ (see Section 3.4 of Melatos et al. 2020 for a detailed discussion of the timing noise model used here). We then estimate the glitch parameters again using LMFIT to fit the post-glitch evolution of $f^*(t_n)$ [this time including the exponential decay term in equation (4.8)] and separately using TEMPONEST to fit the full dataset. The results are shown in Table 4.4. As in Section 4.4.1 we note that the estimates of the glitch parameters differ substantially between the three methods, as expected, because the estimates are conditional on three different signal models. The HMM and TEMPONEST fits are not significantly improved by including a $\Delta \dot{f}_p$ term, so we quote results from analyses which do not include this term.

4.5 | Undetected glitches

4.5.1 | Upper limits

In addition to searching for new glitches, we place upper limits on the magnitude of undetected glitches using synthetic data injections using LIBSTEMPO (Vallisneri, 2020). The injected glitches include only a permanent step in frequency, with no change in frequency derivative or exponential recovery included. We quote 90% frequentist upper limits for each section of data, denoted $\Delta f^{90\%}$. Frequentist upper limits are based on the




Figure 4.6: (Top) Heatmaps of the marginalised posterior distributions of f (left) and \dot{f} (right) for the data section containing the 1999 glitch described in Section 4.4.2. (Bottom) Sequence of *a posteriori* most probable f (left) and \dot{f} (right) states over time, for the same data section. The values of f and \dot{f} are relative to the pre-glitch timing solution, hence they are flat and centred on zero in the pre-glitch region.

| Parameter (units) | HMM fit (via LMFIT) | TEMPONEST fit | Espinoza et al. (2021) |
|-------------------------------------|------------------------------------|------------------------------------|----------------------------------|
| t_g (MJD) | 51424.85 | 51424.85 | 51425.12 ± 0.02 |
| $\Delta\phi$ (turns) | N/A | $9.3_{-9.4}^{+8.7} \times 10^{-4}$ | N/A |
| Δf_p (Hz) | $(3.47 \pm 0.02) \times 10^{-7}$ | $2.7_{-1.1}^{+0.9} \times 10^{-7}$ | $(2.51 \pm 0.05) \times 10^{-7}$ |
| $\Delta \dot{f}_p$ (Hz s $^{-1}$) | $(-0.91 \pm 0.03) \times 10^{-15}$ | $(0.2 \pm 1.7) \times 10^{-14}$ | $(0.8 \pm 0.2) \times 10^{-14}$ |
| $\Delta \ddot{f}_p$ (Hz s $^{-2}$) | N/A | N/A | $(-2.6 \pm 0.2) \pm 10^{-21}$ |
| Δf_d (Hz) | $(0.80 \pm 0.02) \times 10^{-7}$ | $1.6_{-0.9}^{+1.0} \times 10^{-7}$ | $(1.74 \pm 0.05) \times 10^{-7}$ |
| τ_d (days) | 17.5 ± 1.1 | $31.6_{-10.6}^{+17.9}$ | 31 ± 2 |

Table 4.4: Glitch model parameters returned by the HMM (via LMFIT) and TEMPONEST for the small glitch of 1999 described in Section 4.4.2. In the second and third columns the glitch epoch is taken to be immediately after the last pre-glitch ToA, to match the HMM. With TEMPONEST we fit for $\Delta\phi$ instead of t_g . The given ranges are the 90% confidence intervals in the second and fourth columns, and 90% credible intervals in the third column.




Figure 4.7: Frequentist HMM upper limits on the fractional size of undetected glitches $\Delta f^{90\%}/f$ across the 24 years of data analysed here.

expected detection rate over many trials. In particular the 90% frequentist upper limits are the glitch sizes for which we expect to detect a glitch of that size 9 times out of 10, on average. A detailed description of the procedure used to estimate $\Delta f^{90\%}$ values can be found in Appendix B of Dunn et al. (2022b), except that we opt not to inject red timing noise into the synthetic datasets. Experience indicates that although accounting for timing noise is an important part of setting up the HMM glitch detector, the difference between injecting timing noise and not is quite small when setting upper limits, as long as the HMM parameters match what has been used in the search. For example, when setting upper limits on a stretch of data between MJD 48551 and MJD 48719, we find that the value of $\Delta f^{90\%}$ is 1.6×10^{-8} Hz without timing noise injected, and 1.7×10^{-8} Hz with timing noise injected with parameters $A = 10^{-8.8} \text{ yr}^{3/2}$ and $\beta = 8.3$ [see equation (4.9)].

Fig. 4.7 shows the fractional upper limits $\Delta f^{90\%}/f$ obtained over the full dataset. Values on the order of 10^{-9} are typical, with significant variation between sections of the dataset. In the section of data containing the newly reported glitch the fractional upper limit is 1.0×10^{-9} , and the median fractional upper limit across the full dataset is 1.35×10^{-9} . These limits are broadly consistent with the detection limits of Espinoza et al. (2021), which are roughly between $\Delta f/f = 10^{-10}$ and $\Delta f/f = 10^{-8}$, allowing for the dependence on \dot{f} in those limits. It is challenging to construct Fig. 4.7 without the automated and computationally fast algorithm provided by the HMM.

4.5.2 | Cadence of observations

Somewhat counterintuitively, the upper limits derived from the HMM are elevated by the high cadence of the observations. Across a single ToA gap of length z , the uncertainty in frequency is roughly (Melatos et al., 2020)

$$\delta f_1 = (\sigma_{\text{ToA},1}^2 + \sigma_{\text{ToA},2}^2)^{1/2} z^{-1}, \quad (4.11)$$

where $\sigma_{\text{ToA},i}$ are the uncertainties on the two ToAs comprising the gap, measured in turns. Evidently this quantity increases as z decreases, which degrades the sensitivity of the HMM to small glitches. However, we have the freedom to analyse only a subset of the data in order to increase z . Sensitivity also degrades for long z : the timing noise included in the HMM means that jumps in frequency during long ToA gaps may be accommodated as stochastic wandering rather than detected as a glitch. The scale of frequency wandering due to timing noise included in the HMM is given by

$$\delta f_2 = \sigma_{\text{TN}} z^{3/2}, \quad (4.12)$$

where σ_{TN} is a free parameter which controls the strength of the timing noise, defined by equation (4.4). In this work σ_{TN} is on the order of $10^{-19} \text{ Hz s}^{-3/2}$ to $10^{-18} \text{ Hz s}^{-3/2}$. In balancing the low- and high- z behaviour, we seek to minimise the quadrature sum $\sqrt{\delta f_1^2 + \delta f_2^2}$ with respect to z . This typically leads to an optimal cadence of 7–10 days, significantly longer than the near-daily cadence of much of the data analysed here. However, given the data sectioning described in Table 4.1, reducing the data to achieve this optimal cadence frequently yields fewer than 15 ToAs per section, limiting the ability of the HMM to track the evolution of the pulsar. Hence in this work we do not analyse subsets of the data to achieve the optimal cadence.

We emphasise that the considerations presented here are specific to the HMM method, rather than a generic discussion of the effect of observing cadence on glitch sensitivity.

4.6 | Conclusion

In this paper we present the results of an HMM-based search for undetected glitches in the Vela pulsar using the high-cadence dataset from Mount Pleasant Observatory spanning 24 years. We discover one new glitch, with fractional size $\Delta f/f = (8.19 \pm 0.04) \times 10^{-10}$, making it the smallest glitch in Vela which is clearly distinguished from timing noise. We also confidently detect two other small glitches first reported by Espinoza et al. (2021), and re-measure the parameters of these glitches using both the HMM and TEMPO-NEST. The parameter estimates are broadly consistent with one another and with the estimates given by Espinoza et al. (2021), when one considers that they are conditional on three different signal models. We also employ Monte Carlo simulations to place 90% frequentist upper limits on the size of undetected glitches, obtaining a median fractional upper limit of $\Delta f^{90\%}/f = 1.35 \times 10^{-9}$. The latter exercise benefits from the automated and computationally fast nature of the HMM.

These results represent further steps towards a comprehensive understanding of glitch statistics. Such an understanding is crucial to discriminating between falsifiable, microphysics-agnostic models of the glitch mechanism on the basis of long-term glitch statistics, e.g. auto- and cross-correlations between sizes and waiting times (Melatos et al., 2018; Carlin and Melatos, 2019a). It is also important to understanding correlations between glitch activity and pulsar parameters, e.g. the spindown rate or age of the pulsar (Fuentes et al., 2017; Millhouse et al., 2022). Quantifying the completeness of glitch catalogues is also important for analyses which are not directly aimed at investigating glitches, but are nonetheless sensitive to their presence, for example measurements of braking indices of glitching pulsars (Lyne et al., 1996; Espinoza et al., 2017; Lower et al., 2021). For pulsars

of particular interest (e.g. PSR J0534+2200 or PSR J0537–6910) it may well be worth applying similar techniques as those employed here to search for missed glitches.

Acknowledgements

Parts of this research are supported by the Australian Research Council (ARC) Centre of Excellence for Gravitational Wave Discovery (OzGrav) (project number CE170100004) and ARC Discovery Project DP170103625. LD is supported by an Australian Government Research Training Program Scholarship. CME acknowledges support from the ANID FONDECYT grant 1211964. DA acknowledges support from an EPSRC/STFC fellowship (EP/T017325/1). This work was performed in part on the OzSTAR national facility at Swinburne University of Technology. The OzSTAR program receives funding in part from the Astronomy National Collaborative Research Infrastructure Strategy allocation provided by the Australian Government. This work was performed in part on the MASSIVE computing facility, with access provided through the National Computational Merit Allocation Scheme.

Data Availability

The data underlying this article will be shared on reasonable request to the corresponding author.

4.A1 | MJD 48636 event as a candidate in previous searches

The glitch event at MJD 48636 described in Section 4.3 has not been reported until now. However, it did appear as a marginal candidate in the search for glitches presented by Espinoza et al. (2021). In this appendix we briefly discuss the characteristics of this marginal candidate as it appeared in that search, and the reasons it was ultimately not recognised as a genuine glitch at that time.

After the two events which Espinoza et al. (2021) reported as new glitches, the event at MJD 48636 was the next most significant. It was ultimately discarded for the following reasons:

- No $\Delta\dot{f}$ term was detected. This is due to a combination of factors. The automated glitch detection method fits only small stretches of data at once (10 or 20 days), and the event at MJD 48636 is followed by just seven days of data and then a 2-week gap. These factors together are likely to be the reason that $\Delta\dot{f}$ was not detected.
- The detected Δf value was relatively small, and comparable to many other similar events which were classified as noise. No quantitative difference was found between the event at MJD 48636 and these other events.
- In general large data gaps such as the one which follows a few days after this glitch were treated with caution, and strong claims about marginal candidates involving these large gaps were avoided.

5

First results from the UTMOST-NS pulsar timing programme

This chapter has been submitted for publication to Monthly Notices of the Royal Astronomical Society as L. Dunn et al., “First results from the UTMOST-NS pulsar timing programme”.

Abstract

The UTMOST-NS pulsar timing programme operated at the Molonglo Observatory Synthesis Telescope from April 2021 to June 2023, observing 173 pulsars with an average cadence of 50 pulsars per day. An overview of the programme is presented, detailing the hardware, software, and observing strategy. Pulsar timing results are discussed, focusing on timing noise and glitches. It is shown that the scaling of residuals due to timing noise with pulsar parameters and observing timespan is consistent with earlier studies, but that there is a bias towards recovering smaller spectral indices when the observing timespan is increased. Second frequency derivatives are investigated, and it is shown that the uncertainty on $\dot{\nu}$ is sensitive to the frequency cutoff in the timing noise model, varying by three-fold approximately depending on whether Fourier modes with frequency lower than the reciprocal of the observing timespan are included. When low-frequency modes are excluded we measure nine non-zero values of $\dot{\nu}$. If these modes are included then one non-zero $\dot{\nu}$ is measured. An analytic scaling relating anomalous braking indices to timing noise amplitude is also validated. Glitches in the sample are discussed, including three detected by an “online” glitch detection pipeline using a hidden Markov model (HMM). In total 17 glitches are discussed, one of which, in PSR J1902+0615, has not been reported elsewhere. An “offline” glitch search pipeline using the HMM framework is used to search for previously undetected glitches. Systematic upper limits are set on the size of undetected glitches. The mean upper limit is $\Delta\nu^{90\%}/\nu = 5.3 \times 10^{-9}$ at 90% confidence.

5.1 | Introduction

Precision studies of the rotational behaviour of slow (spin period $\gtrsim 0.1$ s) radio pulsars can provide a wealth of information on the extreme environments in and around neutron stars. In particular, rotational irregularities known as “timing noise” and “glitches” can be used to constrain the composition and internal structure of these objects, and the nature of the coupling between the various components (Haskell and Melatos, 2015; Chamel, 2017).

Timing noise appears as a low-level, persistent, time-correlated stochastic wandering in the phase of the pulsar. The details differ from pulsar to pulsar, in some cases being well-modelled as a random walk in the phase, frequency or frequency derivative (Boynton et al., 1972; Cordes and Helfand, 1980), while in other cases it appears that a random walk is unable to account for the observed behaviour (Cordes and Downs, 1985; D’Alessandro et al., 1995; Hobbs et al., 2010). Timing noise is frequently decomposed as a sum of sinusoids with a power-law spectrum, possibly with a corner frequency or a high-frequency cut-off (Lentati et al., 2013; Parthasarathy et al., 2020; Reardon et al., 2023). Phenomenological modelling of this type does not directly address what causes the behaviour in the first place, although a number of mechanisms have been proposed, often involving a superfluid component in the interior of the star (Alpar et al., 1986; Jones, 1990a; Melatos and Link, 2014; Meyers et al., 2021a).

Timing noise is an important confounding factor for performing precision pulsar timing measurements. In pulsar timing this often arises in the context of measuring the second frequency derivative $\ddot{\nu}$, which is of particular physical interest because of its connection with the long-term braking mechanisms at play (Lyne and Graham-Smith, 2012). Values of $\ddot{\nu}$ derived from analyses which are not timing noise-aware are often contaminated (Hobbs et al., 2010; Chukwude and Chidi Odo, 2016; Onuchukwu and Legahara, 2024) – such values of $\ddot{\nu}$ are in some cases themselves used as a measure of timing noise strength (Taylor, 1991; Matsakis et al., 1997). However, under certain circumstances even analyses which do incorporate a timing noise model can fail to fully account for the influence of timing noise on the measured value of $\ddot{\nu}$ (Vargas and Melatos, 2023a; Keith and Nițu, 2023).

Glitches, in contrast to the continuous, low-level variations of timing noise, are step changes in the pulse frequency. They are often associated with a step change in frequency derivative, and an exponential relaxation back towards the pre-glitch trend (Lyne and Graham-Smith, 2012). This phenomenological model of a glitch and its recovery is widely used, but as with timing noise the question of the underlying physical mechanism is far from settled (see e.g. Haskell and Melatos (2015) for a review). Often the superfluid interior is invoked, as significant angular momentum may be transferred from this component to the solid crust, to which the electromagnetic emission is tied (Alpar et al., 1984b; Link and Epstein, 1991; Jones, 1998; Peralta et al., 2006). It has also been suggested that mechanical failure of the crust under strain (a “starquake”) may play an important role in the glitch mechanism (Ruderman, 1976; Link and Epstein, 1996; Middleditch et al., 2006; Kerin and Melatos, 2022).

The mechanisms involved in the glitch phenomenon may be probed both by close study of individual glitches, when sufficiently detailed data are available (Dodson et al., 2002; Palfreyman et al., 2018; Ashton et al., 2019a), and population-level studies investi-

gating both the statistics of glitches in individual pulsars (Melatos et al., 2008; Espinoza et al., 2014; Melatos et al., 2018; Carlin and Melatos, 2019a; Carlin and Melatos, 2019b; Carlin and Melatos, 2020; Fuentes et al., 2019; Ho et al., 2020) as well as the statistics across the pulsar population as a whole (Lyne et al., 2000; Espinoza et al., 2011; Fuentes et al., 2017; Eya et al., 2019; Millhouse et al., 2022). The latter especially are reliant on large-scale observing campaigns (Espinoza et al., 2011; Yu et al., 2013; Jankowski et al., 2019; Basu et al., 2022; Keith et al., 2024) and reliable glitch detection strategies (Espinoza et al., 2014; Yu and Liu, 2017; Melatos et al., 2020; Singha et al., 2021a).

In pursuit of a dataset which could form a platform for studies of these phenomena, the UTMOST-NS pulsar timing programme operated from April 2021 to June 2023, monitoring 173 pulsars with daily to weekly cadence. The UTMOST-NS programme used the refurbished north-south arm of the Molonglo Observatory Synthesis Telescope (MOST), and was undertaken as part of the broader UTMOST project (Bailes et al., 2017; Deller and Flynn, 2020; Day, 2022; Mandlik et al., 2024) which had already revived the east-west arm as a pulsar timing instrument (hereafter referred to as UTMOST-EW) (Jankowski et al., 2019; Lower et al., 2020). In large part the UTMOST-NS programme was a continuation of the UTMOST-EW timing effort, with similar motivation and significant overlap in targets.

The paper begins with a discussion of the observational foundations of the timing programme (Section 5.2), from the details of the telescope configuration to the observing strategy and the radio frequency interference environment. We then set out the pulsar timing framework which is used throughout the rest of this work (Section 5.3), before moving on to investigations of timing noise (Section 5.4), second frequency derivatives (Section 5.5), and glitches (Section 5.6) in the context of the UTMOST-NS dataset. We conclude and look to future work in Section 5.7.

5.2 | Observations

5.2.1 | System overview

The pulsar timing programme is performed using the north-south arm of the Molonglo Observatory Synthesis Telescope (MOST), which was refurbished as part of the UTMOST¹ project (Bailes et al., 2017; Deller and Flynn, 2020; Day, 2022). The refurbished system is referred to as UTMOST-NS. A complete description of the system can be found in Mandlik et al. (2024); here we give a brief overview of its basic characteristics.

The north-south arm of the MOST is composed of two $778\text{ m} \times 12.7\text{ m}$ paraboloid cylindrical reflectors, with a total collecting area of 19800 m^2 . UTMOST-NS consists of 66×1.4 meter long individual array elements called *cassettes*, each containing eight dual-pole patch antennas, which are primarily installed in a “dense core” near the centre of the arm. The total collecting area visible to these cassettes is 1375 m^2 .

The operating frequency range of UTMOST-NS is 810–860 MHz. The typical system equivalent flux density (SEFD) for the timing measurements is 680 Jy, and the typical system temperature is 170 K. The performance of the telescope over the duration of the programme and the estimation of flux densities for pulsars in our sample are discussed

¹UTMOST is not an acronym.

in Appendix 5.A1.

UTMOST-NS operates as a transit instrument with no moving parts — pointing along the meridian is performed electronically via beamformers. The primary beam has an extent of 2.5 degrees in the east-west direction, 12.7 degrees in the north-south direction, and the beam can be pointed in the north-south direction from -90° to approximately $+25^\circ$ in declination. Pulsars are thus typically visible within the half-power points of the primary beam for between 5 and 25 minutes, depending on their declination.

Sky frequencies are down-converted using a 900 MHz signal, which is referenced to a station clock, consisting of a Valon 5007 frequency synthesiser disciplined by a 10 MHz GPS standard.

We primarily use the bright millisecond pulsar PSR J0437–4715 to monitor and correct for small clock jumps, which typically take place after losses of power to the entire system, or other disruptions. Over the course of the timing program, six such corrections are logged, lying in the range $4.2 \mu\text{s}$ to $18 \mu\text{s}$. These corrections are cross-checked for consistency against another millisecond pulsar, J2241–5236. This low dispersion measure (DM) pulsar scintillates and is used as a consistency check on the clock rather than as a primary reference because it was not always detected in daily transits.

For timing purposes, we take the geographical reference position of the telescope to be the same position used for the pulsar timing programme on the east-west arm: a latitude of $-35^\circ 22' 14.5518''$, a longitude of $149^\circ 25' 28.8906''$, and an elevation of 741 meters in the International Terrestrial Reference Frame (2008) coordinates assuming a Geodetic Reference System 1980 ellipsoid. The latitude and longitude were determined in a geodetic survey made by Geoscience Australia in Nov 2012, and the altitude was measured using a consumer GPS device (Garthwaite et al., 2013; Jankowski et al., 2019).

5.2.2 | Target selection and observing strategy

The primary science goal for the UTMOST-NS timing programme is the study of timing irregularities in canonical pulsars². The target selection and observing strategy is focused on high-cadence (typically $\lesssim 5$ d) monitoring of those canonical pulsars which have high enough mean flux densities ($\gtrsim 5$ mJy at 840 MHz, see Appendix 5.A1) to be seen in the time they take to transit the primary beam.

An initial list of approximately 250 candidate pulsars was selected using the mean flux densities in the timing program on the east-west arm of the array (Jankowski et al., 2019). Timing data were acquired for this set of pulsars during commissioning of the upgrade to the north-south arm in 2020 and 2021, and the timing list pared down to 173 pulsars to be monitored regularly based on whether the pulsars were successfully detected in these early observations.

Although the majority of the target pulsars are canonical, we include a handful of millisecond pulsars (MSPs), primarily for the purpose of monitoring the stability of the system clock. These are also particularly useful, because high-cadence monitoring of MSPs serves as a cross-check of timing integrity, when glitches or other step-changes in behaviour are found by our real-time detection system, such as the transient profile event observed in PSR J1737+0747 (Xu et al., 2021; Singha et al., 2021b; Jennings et al., 2022). We also include the radio-loud magnetars PSR J1622–4950 (Levin et al., 2010) and PSR

²That is, slow (spin period $\gtrsim 20$ ms), non-binary, non-magnetar pulsars.

J1809–1943 (Camilo et al., 2006; Levin et al., 2019). Although we are not able to establish timing solutions for either magnetar due to low numbers of detections, PSR J1809–1943 is particularly valuable as a testbed for the Fast Radio Burst (FRB) search pipeline “Adjacent Beam Classifier” (ABC) developed for commensal operation with pulsar timing (Mandlik, 2024).

The observing schedule is constructed semi-automatically. Many targets can be observed daily. In cases where multiple pulsars overlap in right ascension and the available transit time does not permit us to dwell on each for long enough to create a high-quality ToA, we cycle through sets of targets on consecutive days. Broadly speaking, pulsars in regions away from the Galactic plane are mostly observed daily, while pulsars in busy regions of the Galactic plane are observed with a 2- to 6-day cadence. Fig 5.1 shows the locations of the 173 pulsars, and indicates regions of the sky in which daily, and 2- to 6-day cadences are performed by cycling through the available pulsars in transiting regions of sky. The observations are conducted automatically around the clock, based on these schedules, using the HADES observing system. ABC searches for FRBs both during pulsar observations, and while the scheduler is waiting for the next program pulsar to transit the beam (Mandlik et al., 2024).

5.2.3 | RFI environment

As with all radio facilities, contamination of pulsar observations by radio frequency interference (RFI) presented challenges over the course of the timing programme. In particular, the UTMOST-NS operational frequency band (810–860 MHz) overlaps with handset emissions from two Australian mobile phone networks, as well as local sources of narrowband interference (aka ‘birdies’). Very strong signals from 3G and 4G mobile phone base-stations occur at 780–790 MHz and 870–890 MHz, on either side of our operating band, requiring sharp roll-off high- and low-pass filters at the band edges. Due to the site location, which is a flat valley surrounded by hills, and the low local population density, these transmissions rarely rendered an entire observation unusable. Rather, they are well localised in time (typically lasting less than a minute) and frequency (occupying at most 20% of the total observing band). This typical pattern of RFI is cleaned very successfully using XPROF, a tool which uses knowledge of the pulsar profile, flux density and ephemeris to judiciously clean off-pulse regions in sub-integration and frequency channel space.

Figure 5.2 shows an example observation contaminated by cell network signals before and after cleaning. The frequency bands occupied by two cell network transmissions are clearly visible in the pre-cleaning frequency-phase plot (top left), occupying roughly 825–830 MHz and 840–845 MHz. Each transmission occupies only one or two 20 s subintegrations, as shown in the pre-cleaning time-phase plot (top middle). The right-most plots show the fully integrated profiles pre- and post-cleaning. Although the pulse profile is clear in both plots, the pre-cleaning profile shows obvious contamination in the off-pulse region, which is largely removed by the cleaning procedure.

For the pulsars that were regularly timed in the programme, the overall median percentage of data deleted in an observation is 16%. This includes the automatic 10% deleted at the edges of the frequency band, as these regions are often overwhelmed by loud RFI, which interferes with XPROF’s ability to properly clean the data. The per-pulsar medians




Figure 5.1: Sky locations of the 173 pulsars observed in the UTMOST-NS survey, shown in RA in hours and DEC in degrees. The symbol size is proportional to mean flux density. Vertical grey zones show groupings of pulsars based on proximity in RA. A single black symbol in such a zone indicates a pulsar which can be timed daily; green/blue symbols show two pulsars which are timed on alternative days; green/blue/red shows pulsars timed every third day on a three day cycle, up to a maximum of six pulsars timed over six transit days in busy regions of the Galactic plane (particularly between 15 and 20 hours RA, where the plane lies approximately north-south on the sky and there are many pulsars transiting the meridian). The horizontal width of coloured boxes containing each pulsar indicates the integration time (and is on the same scale as the RA axis, i.e. 1 hour in RA is 1 hour of elapsed time).




Figure 5.2: Before (*top row*) and after (*bottom row*) example demonstrating successful RFI cleaning using xPROF for the observation of PSR J1534–5334 taken at 2021-10-08-04:29:08, with the frequency-phase (*left*) and time-phase (*middle*) heatmaps, and the fully integrated pulse profile (*right*). Significant contamination from multiple cell network transmissions (825–830 and 840–845 MHz) is visible in the pre-cleaning plots, and largely absent in the post-cleaning plots.

range between 11% and 24%.

5.2.4 | Timing pipeline

There are several intermediate stages between the data stream acquired at the telescope and the pulse times of arrival which are the primary final data product of the timing programme.

The initial folded archives are produced from tied-array beam data at $10.24 \mu\text{s}$ time resolution using `DSPSR` (van Straten and Bailes, 2011) which performed the coherent dedispersion and folding according to a stored ephemeris. These archives have 512 frequency channels covering the 50 MHz bandwidth and two orthogonal linear polarisations, although only the total intensity has been used in this work.

Once the raw folded archives have been produced, they are passed through the RFI mitigation step – see Section 5.2.3 for a description of the RFI environment at Molonglo and the mitigation procedures.

The RFI-cleaned archives are finally processed into ToAs using `pat`, a component of `PSRCHIVE` (Hotan et al., 2004), to cross-correlate the observed pulse profile against the standard profile and thus derive a pulse arrival time and associated uncertainty. The standard profiles are generated by smoothing high signal-to-noise ratio (S/N) observations of each pulsar. Where available, we use the standard profiles generated for the UTMOST-EW pulsar timing programme (Jankowski et al., 2019; Lower et al., 2020). For pulsars which were not observed on the east-west arm, we obtained standard profiles from observations taken by the MeerKAT telescope (Johnston et al., 2020), or from UTMOST-NS observations. The generated ToAs and uncertainties are stored along with additional metadata (e.g. S/N for the observation, duration of the observation). Only ToAs with S/N greater than 7 are regarded as “good” ToAs to be used in further timing analysis. In total 18360 such ToAs pass this cut out of the final data set of 34575 UTMOST-NS observations. 370 observations which produce a ToA with $S/N > 7$ were judged to be RFI-contaminated and manually excluded, leaving 17990 science-ready ToAs.

As a demonstration of the timing accuracy achieved by the UTMOST-NS system, Figure 5.3 shows the timing residuals for the UTMOST-NS observations of the millisecond pulsar PSR J2241–5236, covering the time period between May 2021 and June 2023 and comprising 120 times of arrival. The mean ToA uncertainty is $4.8 \mu\text{s}$, and the weighted rms of the residuals over the full observation period is $3.2 \mu\text{s}$.

5.3 | Pulsar timing

In this section we introduce the primary pulsar timing methodology used throughout this work. Section 5.3.1 presents the components of the timing model, including both deterministic and stochastic contributions, and Section 5.3.2 discusses the Bayesian model selection procedure used to decide which components should be incorporated into the timing model for a dataset.




Figure 5.3: Timing residuals from UTMOST-NS observations of the millisecond pulsar PSR J2241–5236. The weighted rms is $3.2 \mu\text{s}$.

5.3.1 | Timing framework

We adopt a pulsar timing framework which uses the `ENTERPRISE` (Ellis et al., 2020) pulsar inference package in conjunction with `LIBSTEMPO` (Vallisneri, 2020) and `TEMPO2` (Hobbs et al., 2006). `ENTERPRISE` allows the user to construct a model for the pulsar ToAs which incorporates both deterministic and stochastic processes, i.e. $\mathbf{t} = \mathbf{t}_{\text{det}} + \mathbf{t}_{\text{sto}}$, where \mathbf{t} is the vector of ToAs. The deterministic contribution is generated by the standard phase model provided by `TEMPO2` (Edwards et al., 2006), and contains both the intrinsic deterministic evolution of the pulsar, expressed as a low-order Taylor series in the time coordinate of the pulsar’s reference frame, as well as the kinematic Doppler contribution associated with the transformation between the pulsar reference frame and the reference frame of the telescope. While there are in principle a large number of parameters which determine the form of \mathbf{t}_{det} , in this work we primarily restrict our attention to the pulsar’s intrinsic rotational frequency and its first two derivatives, ν , $\dot{\nu}$, and $\ddot{\nu}$, as well as the pulsar’s position on the sky, expressed in terms of right ascension α and declination δ . In some cases further parameters are necessary to obtain a good fit, e.g. binary orbital elements and proper motions, but we leave these fixed at the values recorded in the first UTMOST data release (Lower et al., 2020) where applicable, or the ATNF pulsar catalogue (Manchester et al., 2005). In cases where a clear glitch has been detected from visual inspection of the residuals we include a corresponding glitch term in our phase model, viz.

$$\begin{aligned} \phi_g(t) = & \left[\Delta\phi_p + \Delta\nu_p(t - t_g) + \frac{1}{2}\Delta\dot{\nu}_p(t - t_g)^2 \right. \\ & \left. - \tau_d\Delta\nu_d e^{-(t-t_g)/\tau_d} \right] \Theta(t - t_g). \end{aligned} \quad (5.1)$$

Here $\Delta\phi_p$, $\Delta\nu_p$ and $\Delta\dot{\nu}_p$ denote permanent step changes in ϕ , ν , and $\dot{\nu}$ respectively. The exponential term represents a decaying step in frequency of magnitude $\Delta\nu_d$ which relaxes on a timescales τ_d . The factor $\Theta(t - t_g)$ is the Heaviside step function, and t_g indicates the glitch epoch. All parameters except t_g are allowed to vary when inferring the glitch properties.

There are two stochastic processes of interest. The first is white noise in the residuals due to ToA measurement uncertainty σ_{ToA} , calculated by `pat`, and two additional parameters EFAC and EQUAD, which account for misestimation of σ_{ToA} due to e.g. non-ideal profile templates or intrinsic profile variation (Liu et al., 2011; Shannon et al., 2014), viz.

$$\sigma_{\text{out}}^2 = (\text{EFAC})^2 \sigma_{\text{ToA}}^2 + (\text{EQUAD})^2. \quad (5.2)$$

The second stochastic process is a long-term “wandering” in the residuals, commonly known as timing noise. The contribution to \mathbf{t} from the latter process is assumed to have a power spectral density (PSD) of the form

$$P(f) = \frac{A_{\text{red}}^2}{12\pi^2 f_{\text{yr}}^3} \left(\frac{f}{f_{\text{yr}}} \right)^{-\beta}, \quad (5.3)$$

and is modelled in `ENTERPRISE` as a set of harmonically related sinusoids. Note that in this parametrisation A_{red} is dimensionless. By default the sinusoids have their lowest frequency at $1/T_{\text{span}}$ and extend up to $30/T_{\text{span}}$ in increments of $1/T_{\text{span}}$. However, this is easily modified, and we also consider a “long-period” timing noise model where the lowest frequency is instead at $1/(2T_{\text{span}})$, but still extending up to $30/T_{\text{span}}$ in increments of $1/(2T_{\text{span}})$, so this model has twice as many sinusoids overall. Although it is possible to fit for the amplitude of each harmonic separately and then combine these into an estimate of the form of the PSD, here we opt instead to constrain the PSD to take the form of equation (5.3) and estimate A_{red} and β . In this approach, A_{red} and β play the role of hyperparameters: the prior distributions on the amplitudes of each harmonic are fixed by $P(f)$, and subsequently marginalised over (Lentati et al., 2013). We do not include models for chromatic stochastic processes, such as variations in the DM or scattering (Goncharov et al., 2021). Given the small fractional bandwidth of our observations, these effects are covariant with achromatic red noise (Lower et al., 2020).

We adopt a Bayesian approach to inferring the various parameters of the different processes in the pulsar timing models. `ENTERPRISE` constructs an appropriate likelihood function for a model \mathcal{M} , $\mathcal{L}(\mathbf{t} | \boldsymbol{\theta}, \mathcal{M})$, where $\boldsymbol{\theta}$ is a vector of parameter values. Given this likelihood along with a prior distribution on the parameters of interest, $\text{Pr}(\boldsymbol{\theta})$, we use the `MULTINEST` (Feroz et al., 2009) nested sampling package to explore the parameter space and estimate the posterior distribution of $\boldsymbol{\theta}$,

$$\text{Pr}(\boldsymbol{\theta} | \mathbf{t}, \mathcal{M}) = \frac{\mathcal{L}(\mathbf{t} | \boldsymbol{\theta}, \mathcal{M}) \text{Pr}(\boldsymbol{\theta})}{\text{Pr}(\mathbf{t} | \mathcal{M})} \quad (5.4)$$

as well as the model evidence,

$$\text{Pr}(\mathbf{t} | \mathcal{M}) = \int d\boldsymbol{\theta} \mathcal{L}(\mathbf{t} | \boldsymbol{\theta}, \mathcal{M}) \text{Pr}(\boldsymbol{\theta}). \quad (5.5)$$

The model evidence is also denoted by \mathcal{Z} .

| Parameter | Prior range |
|--|---|
| $\alpha, \delta, \nu, \dot{\nu}$ | $p^* \pm 5000\Delta_p$ |
| $\dot{\nu}$ | $\pm 5000\sigma_{\text{rms}}T_{\text{span}}^{-3}$ |
| $\Delta\phi, \Delta\dot{\nu}_p$ | $p^* \pm 5000\Delta_p$ |
| $\log_{10}(\Delta\nu_p/\text{Hz}), \log_{10}(\Delta\nu_d/\text{Hz})$ | $[-10, -4]$ |
| $\log_{10}(\tau_d/\text{d})$ | $[-1, 3]$ |
| $\log_{10}(\text{EFAC})$ | $[-1, 0.7]$ |
| $\log_{10}(\text{EQUAD}/\text{s})$ | $[-8, -2]$ |
| $\log_{10}(A_{\text{red}}/\text{yr}^{3/2})$ | $[-14, -7]$ |
| β | $[2, 10]$ |

Table 5.1: Prior ranges used in ENTERPRISE analyses described in Section 5.3. All prior distributions are (log-)uniform. For a timing model parameter p , the best-fit value of the parameter returned by an initial TEMPO2 fit is denoted by p^* , and the associated uncertainty is denoted by Δ_p . The rms of the phase residuals from the TEMPO2 fit is denoted by σ_{rms} .

5.3.2 | Parameter inference and model selection

To go from the pulsar timing framework described in the previous section to inference results for the pulsars in our sample, we must first decide on what data are to be incorporated into the analyses. For those pulsars which were timed as part of both the UTMOST-EW and UTMOST-NS timing programmes, we combine the two datasets to maximise the timing baseline. This includes a majority of the UTMOST-NS pulsars, with 93% (150) having previous data from the UTMOST-EW programme, and 83% (134) with a total timing baseline of more than 7 years. Those pulsars for which only UTMOST-NS data exist have timing baselines between approximately 6 months and 2 years.

The values of the timing model parameters are informed by an initial least-squares fit performed using TEMPO2, which for each timing model parameter p returns a best-fit value p^* and an associated error estimate Δ_p . The priors on the timing model parameters are then taken to be uniform on $p^* \pm 5000\Delta_p$, with the following exceptions. When performing the initial least-squares fit we do not fit for $\dot{\nu}$, and the prior on $\dot{\nu}$ is taken to be $\pm 5000\sigma_{\text{rms}}T_{\text{span}}^{-3}$. If a glitch is included in the fit, then the initial least-squares fit is performed with no decay term included. The parameters $\Delta\phi$ and $\Delta\dot{\nu}_p$ are treated in the same way as other timing model parameters. However, the permanent and decaying frequency increment are taken to have a log-uniform distribution on $[10^{-10} \text{ Hz}, 10^{-4} \text{ Hz}]$, while the decay timescale is taken to have a log-uniform distribution on $[10^{-1} \text{ d}, 10^3 \text{ d}]$. The log-uniform distributions on the frequency increments reflect the fact that *a priori* their individual orders of magnitude are not well-constrained, even though their sum is reasonably well-measured by the initial least-squares fit. See Section 5.6.5 for further discussion on the treatment of the decaying component in the glitch parameter estimation.

We do not necessarily expect all of the processes described in Section 5.3.1 to be present in every dataset. A minimal model includes deterministic residuals due to timing model parameter errors and stochastic residuals due to ToA measurement error, but not every pulsar displays detectable levels of timing noise, and many pulsars have no detectable

second frequency derivative. To assess the presence of these additional processes in each dataset we perform model selection by computing the model evidence \mathcal{Z} for six models (itemized) using MULTINEST, and the Bayes factors $\mathcal{B}_{12} = \mathcal{Z}_1/\mathcal{Z}_2$ indicating the pairwise favouring of one model over another (assuming all models are *a priori* equally likely). For each pulsar we select between six models:

- NIL: This is the minimal model, which only includes contributions from a timing model with $\ddot{\nu} = 0$ and white measurement noise.
- TN: This includes everything from NIL, plus a red noise component parameterised by A_{red} and β [see equation (5.3)], with low-frequency cutoff at $1/T_{\text{span}}$.
- TNLONG: The same as TN, except that the low-frequency cutoff is instead at $1/(2T_{\text{span}})$.
- F2: This includes everything from NIL, plus a deterministic component due to the presence of a $\ddot{\nu}$ term.
- TNF2: This is the maximal model, including everything from NIL, plus the red noise component as in TN and the second frequency derivative as in F2.
- TNLONGF2: The same as TNF2, except that the low-frequency cutoff is instead at $1/(2T_{\text{span}})$.

In choosing a preferred model for each pulsar to be used in further analyses, we do not simply take the model with the highest \mathcal{Z} . Instead, for each additional “component” over the NIL model (i.e. either timing noise or a non-zero $\ddot{\nu}$) we require a model containing that component to be preferred by $\ln \mathcal{B} > 5$ over every model which does not contain that component. For example, in the case of PSR J1709–4429, TNLONG is favored over NIL by $\ln \mathcal{B} = 4.3 \times 10^3$, and TNF2 is favored over TNLONG by $\ln \mathcal{B} = 0.62$. The data clearly favour a model which includes timing noise, but at best weakly prefer a model which includes a $\ddot{\nu}$ term in addition to timing noise. In this case (ignoring the other models for simplicity) we select TNLONG as the preferred model, as there is only marginal evidence for the $\ddot{\nu}$ term. We do not impose the requirement $\ln \mathcal{B} > 5$ when choosing between TN and TNLONG models – if both are preferred over the non-timing noise models with $\ln \mathcal{B} > 5$ we simply take whichever has the higher \mathcal{Z} as the model going forward. In cases where specific values of the model parameters are needed, we take the mean posterior values for the preferred model.

5.4 | Timing noise

A large fraction of the pulsars in our sample exhibit detectable levels of timing noise, and in this section we investigate its properties across the sample. In total there are 118 pulsars in our sample for which a model including a timing noise component is preferred by a log Bayes factor > 5 over any model which does not include timing noise, and these are the pulsars which we take as having “significant” timing noise.

In order to compare timing noise between datasets, it is useful to adopt some standard measure of timing noise strength which gives an indication of the contribution from the modelled timing noise to the timing residuals. Here we follow Parthasarathy et al. (2019)

and Lower et al. (2020) and adopt the following figure of merit, based on the values of the timing noise parameters returned by ENTERPRISE [see equation (5.3)] and the low-frequency cutoff in the timing noise model, f_{low} :

$$\sigma_{\text{RN}}^2 = \int_{f_{\text{low}}}^{\infty} P(f) df \quad (5.6)$$

$$= f_{\text{yr}}^{-3+\beta} \frac{A_{\text{red}}^2}{12\pi^2} \frac{f_{\text{low}}^{-(\beta-1)}}{\beta-1}. \quad (5.7)$$

We set $1/T_{\text{span}}$ for TN models, $1/2T_{\text{span}}$ for TNLONG models.

5.4.1 | Correlation between timing noise activity and spin parameters

A number of authors have investigated how timing noise strength correlates with pulsar spin parameters [e.g. Cordes and Helfand (1980), Dewey and Cordes (1989), Urama et al. (2006), Shannon and Cordes (2010), Hobbs et al. (2010), Melatos and Link (2014), Parthasarathy et al. (2019), and Lower et al. (2020)], generally finding that there is a significant correlation of timing noise activity with $\dot{\nu}$, for a variety of activity measures and a variety of methodologies for assessing the correlations. Here, following Dewey and Cordes (1989), we look to estimate σ_{RN} based on the spin frequency and frequency derivative, i.e. we construct

$$\chi_{\text{RN}} = C\nu^a |\dot{\nu}|^b f_{\text{low}}^{-\gamma} \quad (5.8)$$

and aim to estimate C , a , b , and γ such that χ_{RN} accurately estimates σ_{RN} (χ_{RN} has units of seconds, as does σ_{RN}).

As in Shannon and Cordes (2010) we assume that the residuals between χ_{RN} and σ_{RN} are log-normally distributed and hence adopt the likelihood

$$\mathcal{L} = \prod_{i=1}^N \frac{1}{\sqrt{2\pi\epsilon}} \exp \left\{ -\frac{[\log_{10}(\sigma_{\text{RN},i}) - \log_{10}(\chi_{\text{RN},i})]^2}{2\epsilon^2} \right\}, \quad (5.9)$$

where the product is over all N pulsars with significant timing noise, indexed by i , and ϵ is an additional parameter which describes the scatter in the relation between σ_{RN} and χ_{RN} . The scatter in the relation may arise from random variation in the timing noise realisations, or from variations in other parameters which are not included in χ_{RN} [e.g. neutron star composition and temperature (Alpar et al., 1986; Melatos and Link, 2014)], or it may arise from a breakdown in the simplifying assumption that the values of C , a , b , and γ are universal across the population. For this investigation we do not treat these potentially complex issues, and simply absorb all of these sources of variation into ϵ , except that the question of variation in γ is discussed at the end of this section.

We take (log-)uniform priors, with a , b , γ , $\log_{10} C \in [-10, 10]$ and $\epsilon \in [0, 10]$. We use BILBY (Ashton et al., 2019b) in conjunction with the DYNESTY (Speagle, 2020) nested sampling package to obtain posterior distributions on all parameters – the parameter estimates are summarised in the first column of Table 5.2.

The estimated values of the a , b , and γ parameters agree broadly with the earlier studies of Shannon and Cordes (2010) and Lower et al. (2020) (although we note that our

Table 5.2: Estimates of the parameters determining χ_{RN} [equation (5.8)], as discussed in Section 5.4.1. In this work we consider two cases, one with $\chi_{\text{RN}} \propto f_{\text{low}}^{-\gamma}$ (second column), where γ is a free parameter to be inferred, and one with $\chi_{\text{RN}} \propto f_{\text{low}}^{-(\beta-1)/2}$ (third column). The fourth and fifth columns present the equivalent results from Shannon and Cordes (2010) and Lower et al. (2020), respectively. Error bars are 95% credible intervals for the quoted results from this work and from Lower et al. (2020), and 2- σ confidence intervals for the results of Shannon and Cordes (2010).

| Parameter | This work ($f_{\text{low}}^{-\gamma}$) | This work [$f_{\text{low}}^{-(\beta-1)/2}$] | SC10 | Lo+20 |
|---------------|--|---|-----------------|-------------------------|
| $\log_{10} C$ | -0.4 ± 1.8 | $-4.4^{+2.2}_{-2.1}$ | 2.0 ± 0.4 | $3.7^{+2.4}_{-2.7}$ |
| a | $-0.92^{+0.31}_{-0.29}$ | $-0.17^{+0.35}_{-0.38}$ | -0.9 ± 0.2 | $-0.84^{+0.47}_{-0.49}$ |
| b | $0.80^{+0.12}_{-0.13}$ | 0.54 ± 0.15 | 1.00 ± 0.05 | $0.97^{+0.16}_{-0.19}$ |
| γ | $1.8^{+0.6}_{-0.8}$ | – | 1.9 ± 0.2 | 1.0 ± 1.2 |
| ϵ | 0.69 ± 0.09 | 0.82 ± 0.10 | 0.69 ± 0.04 | $0.64^{+0.11}_{-0.16}$ |

datasets are not independent of Lower et al. 2020, as we include a subset of the UTMOST-EW data which were used in that analysis). As in those cases, we find $a \approx -b$ within the quoted uncertainties, which are comparable; we find roughly ± 0.3 for a and ± 0.1 for b , noting that Shannon and Cordes 2010 quote 2- σ confidence intervals while we quote 95% credible intervals, as do Lower et al. (2020). This suggests that the characteristic age $\tau_c \sim \nu |\dot{\nu}|^{-1}$ is a good proxy for σ_{RN} , at least for canonical pulsars, which make up a majority of this sample. It has been noted that τ_c is also a good proxy for a pulsar’s glitch rate, which may or may not involve related physical processes. Millhouse et al. (2022) showed that a model in which glitch rate depends on τ_c is preferred over a model where the glitch depends on a generic combination $\nu^a \dot{\nu}^b$ with $b \neq a$, or ν^a or $\dot{\nu}^b$ alone. We also find agreement in the value of γ , although there is more variation in both the point estimate and the uncertainties. Both this work and Shannon and Cordes (2010) estimate $\gamma \approx 1.9$, but Shannon and Cordes (2010) report uncertainties of only 0.2 compared to our uncertainty of approximately 0.7. Lower et al. (2020) estimate $\gamma = 1.0 \pm 1.2$, consistent within uncertainty with the other values but less constrained. The degree of scatter is also similar among all three analyses, with $\epsilon \approx 0.7$ in all cases.

For simplicity we assume so far that γ takes one value across the population, but the form of σ_{RN} [equation (5.7)] suggests that if A_{red} is constant over the observing span we should expect $\gamma = (\beta - 1)/2$. We therefore obtain a second set of estimates for C , a , b and ϵ , using the same functional form as equation (5.8) but with $\chi_{\text{RN}} \propto f_{\text{low}}^{-(\beta-1)/2}$, where β is taken to be the mean posterior value from the favoured model obtained via analysis with ENTERPRISE. The results are given in the second column of Table 5.2. The a and b values are significantly different from the $f_{\text{low}}^{-\gamma}$ fit, and are only marginally consistent with $a = -b$. The scatter is also higher for $\chi_{\text{RN}} \propto f_{\text{low}}^{-(\beta-1)/2}$ as indicated by the ϵ values, and the $f_{\text{low}}^{-\gamma}$ model is favoured over the $f_{\text{low}}^{-(\beta-1)/2}$ model by a log Bayes factor of 16. One possible explanation for the poor performance of the $f_{\text{low}}^{-(\beta-1)/2}$ model, explored in Section 5.4.2, is that the measurements of A_{red} and β are not constant as T_{span} increases. In that scenario, the logic which motivates the model no longer applies: the estimated values of A_{red} and β are themselves functions of T_{span} , and the dependence of σ_{RN}^2 on $T_{\text{span}} \sim f_{\text{low}}^{-1}$

is no longer the straightforward $f_{\text{low}}^{-(\beta-1)}$ factor in equation (5.7).

5.4.2 | Consistency of UTMOST-EW and -NS timing noise parameter estimates

In the previous section, we noted that our results are not independent of Lower et al. (2020), as our dataset includes a subset of the full UTMOST-EW dataset analysed in that work, augmented with the addition of the UTMOST-NS data which extends the timespan covered. These two datasets are analysed with similar, but not identical, methods – Lower et al. (2020) used `TEMPONEST` rather than `ENTERPRISE` to perform Bayesian model selection and parameter estimation. In this section we investigate the consistency of the estimates of timing noise parameters with respect to both timing baseline and methodology.

There is no obvious reason why the estimates of the timing noise PSD should differ significantly when `ENTERPRISE` is used instead of `TEMPONEST`, or when additional years of data are added. However, in the latter case it is plausible that the PSD deviates from power-law behaviour as the frequency decreases, which would alter the inferred parameters when fitting to a pure power-law model as more data (and hence lower frequencies) are incorporated (Melatos and Link, 2014; Caballero et al., 2016; Parthasarathy et al., 2019; Goncharov et al., 2020; Antonelli et al., 2023).

The top panel of Figure 5.4 shows the difference in $\log_{10} A_{\text{red}}$ and β between the estimates of this work, using `ENTERPRISE` and the combined EW+NS dataset, and those of Lower et al. (2020), using `TEMPONEST` and the UTMOST-EW data alone. In this comparison we include the pulsars which are common to both analyses and for which the TN model is preferred in the combined EW+NS analysis over the NIL model by $\ln \mathcal{B} > 5$. We also impose a condition for all comparisons in this section, that A_{red} must be greater than 10^{-11} in both analyses. This excludes low-amplitude timing noise for which the PSD parameters are typically not well-constrained [cf. Figure 4 of Vargas and Melatos (2023a)]. For the purposes of these comparisons we do not use the `TNLONG` model. We note that some of the spectral indices reported by Lower et al. (2020) are as large as 20. Here we restrict our attention to pulsars from the Lower et al. (2020) sample which have $2 < \beta < 10$, matching the prior used on β in the present work. For notational convenience we write the differences between the two sets of parameter estimates as $\Delta \log_{10} A_{\text{red}}$ and $\Delta \beta$ where it is clear from context which two sets of parameter estimates we are differencing. For example, in the comparison between the UTMOST-EW+NS `ENTERPRISE` analyses and the UTMOST-EW-only `TEMPONEST` analyses by Lower et al. (2020) we take the differences

$$\Delta \log_{10} A_{\text{red}} = \log_{10} A_{\text{red}}^{\text{EW+NS, ENT}} - \log_{10} A_{\text{red}}^{\text{EW, TEM}}, \quad (5.10)$$


$$\Delta \beta = \beta^{\text{EW+NS, ENT}} - \beta^{\text{EW, TEM}}, \quad (5.11)$$

where we abbreviate the software used as ENT (`ENTERPRISE`) and TEM (`TEMPONEST`).

The mean of the joint distribution of $(\Delta \log_{10} A_{\text{red}}, \Delta \beta)$, indicated by the star in the top panel of Figure 5.4, is $(0.01 \pm 0.04, -0.14 \pm 0.16)$ where the uncertainties are the standard errors of the mean. While there appears to be some scatter between the results obtained in the two sets of analyses, there is no evidence for a systematic offset between them.

[EW+NS (ENTERPRISE)] – [EW (TEMPONEST)]

[EW+NS (ENTERPRISE)] – [EW (ENTERPRISE)]



[EW (ENTERPRISE)] – [EW (TEMPONEST)]




Figure 5.4: Distributions of the differences in recovered timing noise parameters A_{red} and β for various combinations of methodology and dataset, as discussed in Section 5.4.2. The red stars in each panel indicate the locations of the means of the joint distributions of differences — see Table 5.3 for the numerical values of these means. When comparing the EW+NS ENTERPRISE analysis to the EW-only TEMPONEST analysis of Lower et al. (2020) (top left panel), there is no systematic shift in the recovered PSD parameters. When comparing the EW+NS ENTERPRISE analysis to EW-only ENTERPRISE analysis (top right panel), there is a tendency to recover smaller β , but no significant shift in $\log_{10} A_{\text{red}}$. The mean of the joint distribution does not differ significantly from zero when the two analyses of the EW-only data are compared (bottom panel).

Table 5.3: Mean differences in recovered timing noise parameters, $\langle \Delta \log_{10} A_{\text{red}} \rangle$ and $\langle \Delta \beta \rangle$, for three pairings of dataset and method. The dataset is either the combined UTMOST-EW and UTMOST-NS dataset (EW+NS) or the UTMOST-EW data alone (EW). The method is either ENTERPRISE (ENT) or TEMPONEST (TEM). The (EW, TEM) timing noise parameters are taken from the analysis of Lower et al. (2020). The complete distributions of timing noise parameter differences are shown in Figure 5.4. The quoted uncertainty on each mean difference is the standard error of the mean.

| Datasets and methods | $\langle \Delta \log_{10} A_{\text{red}} \rangle$ | $\langle \Delta \beta \rangle$ |
|----------------------------|---|--------------------------------|
| (EW + NS, ENT) – (EW, TEM) | 0.01 ± 0.04 | -0.14 ± 0.16 |
| (EW + NS, ENT) – (EW, ENT) | 0.07 ± 0.04 | -0.57 ± 0.14 |
| (EW, ENT) – (EW, TEM) | 0.01 ± 0.03 | 0.25 ± 0.11 |

However, the estimates in the top panel of Figure 5.4 were obtained using a different method and a different timespan. In order to try and disentangle the effects of these two variables, we re-analyse the UTMOST-EW data using ENTERPRISE. The middle panel shows the difference between the ENTERPRISE results for the combined EW+NS dataset and the ENTERPRISE results for the EW data alone. The distribution of $(\Delta \log_{10} A_{\text{red}}, \Delta \beta)$ shows a significant offset in $\Delta \beta$, with mean $(0.07 \pm 0.04, -0.57 \pm 0.14)$, i.e. there is a tendency to recover somewhat shallower PSDs when the dataset is extended. The bottom panel of Figure 5.4 shows the difference in recovered parameters between the ENTERPRISE re-analysis of the UTMOST-EW data and the TEMPONEST values for the same data. In this case the distribution of $(\Delta \log_{10} A_{\text{red}}, \Delta \beta)$ has a mean marginally different from zero, $(0.01 \pm 0.03, 0.25 \pm 0.11)$. The results of all three pairings discussed are summarised in Table 5.3.

The results obtained in this section suggest that increasing the length of the dataset can introduce a systematic shift in the recovered slope of the timing noise PSD, with shallower PSDs preferred for longer datasets. However, it is unclear whether this effect is due to the character of the timing noise, i.e. deviation from pure power-law behaviour, or whether it is an artifact of the method of timing noise analysis employed in this work. A detailed study of the underlying cause of the variation in the recovered timing noise parameters under changes in method and dataset is beyond the scope of this paper.

If we wish to take A_{red} and β as meaningful parameters in their own right, not only when combined into a figure of merit like σ_{RN} , the variation above presents a potential problem. For example, in Section 5.4.1 we investigate whether using β to fix the scaling of χ_{RN} with observing timespan leads to a tighter correlation between χ_{RN} and σ_{RN} . The logic behind this investigation relies on the estimated values of A_{red} and β being independent of T_{span} . If this is not the case then the scaling implied by equation (5.7) of $\sigma_{\text{RN}}^2 \sim f_{\text{low}}^{-(\beta-1)}$ does not apply. In light of the variation in β with observing timespan observed in this section, it is unsurprising that enforcing the latter scaling does not lead to a tighter correlation between σ_{RN} and χ_{RN} .

5.5 | Second frequency derivatives

Timing noise and the influence of the second frequency derivative are both low-frequency effects, in the sense that their contributions to timing residuals are concentrated in Fourier modes with frequency comparable to T_{span}^{-1} . This can make it difficult to disentangle the influence of the two contributions.

The secular value of $\dot{\nu}$ is of particular physical interest because of its connection to the long-term spin-down of the pulsar. For many spin-down mechanisms the expected frequency evolution has the form

$$\dot{\nu} = K\nu^{n_{\text{pl}}}, \quad (5.12)$$

where, n_{pl} is referred to as the “braking index” and may be estimated (as long as timing noise is low enough; see below) by $n_{\text{pl}} \approx n$ assuming K is a constant, with

$$n = \frac{\nu\ddot{\nu}}{\dot{\nu}^2}. \quad (5.13)$$

The value of n_{pl} differs between physical mechanisms (Blandford and Romani, 1988; Vargas and Melatos, 2023a; Abolmasov et al., 2024), but is typically in the range of $1 \lesssim n_{\text{pl}} \lesssim 7$, e.g. spin-down due to electromagnetic dipole radiation has $n_{\text{pl}} = 3$, spin-down due to gravitational wave emission from a mass quadrupole has $n_{\text{pl}} = 5$, and spin-down due to gravitational wave emission from a mass current quadrupole has $n_{\text{pl}} = 7$ (Lyne and Graham-Smith, 2012; Riles, 2023).

In our sample, only one pulsar is found to favour a model with $\dot{\nu} \neq 0$ and $\ln \mathcal{B} > 5$: PSR J0534+2200 (the Crab pulsar), which has $\dot{\nu} = 1.16(1) \times 10^{-20} \text{ Hz s}^{-2}$ [$n = 2.54(3)$]. The log Bayes factor of the best model for PSR J0534+2200 (TNF2) over the best model which does not include a second frequency derivative (TNLONG) is 61. This is a confident detection, and agrees within error with the interglitch braking index reported by Lyne et al. (2015) of 2.519(2).

If one fits for $\ddot{\nu}$ in a pulsar with significant timing noise, one is liable to find large values of $\ddot{\nu}$, in the sense that the implied braking index is orders of magnitude different from the “canonical” range of $1 \lesssim n \lesssim 7$ and can take either sign. It is well understood that such values are unreliable as indicators of the long-term behaviour, as the residuals due to timing noise can be absorbed into the $\ddot{\nu}$ term (Hobbs et al., 2004; Chukwude and Chidi Odo, 2016). However, with the advent of Bayesian techniques which simultaneously fit for the residual contributions of the timing noise and the $\ddot{\nu}$ term, a number of “anomalous” braking indices with large magnitudes have been reported, with values ranging from -9.6×10^4 to 2.9×10^3 (Lower et al., 2020; Parthasarathy et al., 2020).

Recently it has been pointed out that despite the more sophisticated handling of the mixing between timing noise and $\ddot{\nu}$ in these Bayesian analyses, reported values of n must still be treated with caution in some cases (Vargas and Melatos, 2023a; Keith and Nițu, 2023). In both of these studies it was shown that under certain conditions, the phenomenological timing noise model of equation (5.3), particularly with a low-frequency cut-off at $1/T_{\text{span}}$ (which is the default setting in both ENTERPRISE and TEMPONEST), can lead to erroneous characterisation of the $\ddot{\nu}$ term. In the remainder of this section we discuss the interaction between timing noise and the estimation of uncertainty on $\ddot{\nu}$ (Section 5.5.1), the prospect of measuring n_{pl} for pulsars in our sample (Section 5.5.2),

and the effect of timing noise on the detection of $\ddot{\nu}$ (Section 5.5.3) in the context of our dataset.

5.5.1 | Timing noise and estimated $\dot{\nu}$ uncertainties

We can estimate the expected spread in $\dot{\nu}$ measurements due to the influence of timing noise and compare this to the $\dot{\nu}$ uncertainties returned by ENTERPRISE in the TNF2 and TNLONGF2 models (Vargas and Melatos, 2023a). Timing noise causes the observed $\dot{\nu}$ to change from realisation to realisation, but only one realisation is observed in an astronomical observation, and it is impossible to determine independently whether the observed realisation is typical or rare within the ensemble of physically permitted realisations.

In the first instance, we take the subset of pulsars for which the timing noise parameter estimation for both the TNF2 and TNLONGF2 models returns a spectral index satisfying $|\beta - 6| < 1$. This condition on β suggests that the timing noise may be approximately modeled by a random walk in $\dot{\nu}$, i.e.³

$$\frac{d\dot{\nu}}{dt} = \dot{\nu}_0 + \xi(t), \quad (5.14)$$

where $\dot{\nu}_0$ is an intrinsic constant second frequency derivative and $\xi(t)$ is a white noise term satisfying

$$\langle \xi(t) \rangle = 0, \quad (5.15)$$

$$\langle \xi(t)\xi(t') \rangle = \sigma_{\dot{\nu}}^2 \delta(t - t'). \quad (5.16)$$

This model corresponds to a particular limit of the model used by Vargas and Melatos (2023a). There, the stochastic driving term is in $d\ddot{\nu}/dt$, and $\ddot{\nu}$ and $\dot{\nu}$ revert to their means on timescales $\gamma_{\ddot{\nu}}^{-1}$ and $\gamma_{\dot{\nu}}^{-1}$ respectively [see equations (2)–(5) of the latter reference]. The model given by equations (5.14)–(5.16) here corresponds to the limit $\gamma_{\dot{\nu}}^{-1} \gg T_{\text{span}}$ and $\gamma_{\ddot{\nu}}^{-1} \lesssim T_{\text{cad}}$ where T_{cad} is the observing cadence; cf. equations (A15) and (A16) of Vargas and Melatos (2023a).

The PSD of the phase, $P_{\phi}(f)$, may be estimated by the Wiener-Kinchin theorem, which connects the autocorrelation function with the PSD via a Fourier transform:

$$P_{\phi}(f) = 2\sigma_{\dot{\nu}}^2(2\pi f)^{-6}. \quad (5.17)$$

We then wish to compare $P_{\phi}(f)$ to $P(f)$ inferred from ENTERPRISE which is parametrised as in equation (5.3), and hence obtain an expression for $\sigma_{\dot{\nu}}$ in terms of A_{red} and β . In order to render $P_{\phi}(f)$ and $P(f)$ comparable, there are two factors which must be taken into account. First, $P_{\phi}(f)$ is the PSD of the residuals in units of cycles, while $P(f)$ is the PSD of the residuals in units of seconds. Thus $P(f)$ should be multiplied by a factor of ν^2 in order to match $P_{\phi}(f)$. Second, $P_{\phi}(f)$ has a spectral index of exactly -6 , whereas $P(f)$ generally does not. For the purposes of this comparison we therefore replace $P(f)$ with a new PSD $P^{\beta=6}(f)$ parametrised in the same way as $P(f)$, but with β fixed to 6 and A_{red}

³For simplicity we do not include a mean-reverting term in the equation of motion, on the assumption that the mean-reversion timescale for $\dot{\nu}$ is significantly longer than the observing timespan (Vargas and Melatos, 2023a).

adjusted so that the values of σ_{RN} [equation (5.7)] for $P(f)$ and $P^{\beta=6}(f)$ coincide. This adjusted A_{red} value is denoted by $A_{\text{red}}^{\beta=6}$, and is given by

$$A_{\text{red}}^{\beta=6} = \sqrt{\frac{5}{\beta-1}} A_{\text{red}} \left(\frac{f_{\text{low}}}{f_{\text{yr}}} \right)^{-(\beta-6)/2}. \quad (5.18)$$

Equating $P_{\phi}(f)$ and $\nu^2 P^{\beta=6}(f)$ and isolating $\sigma_{\dot{\nu}}$ gives

$$\sigma_{\dot{\nu}} = \sqrt{\frac{f_{\text{yr}}^3 \nu^2 (2\pi)^6}{24\pi^2}} A_{\text{red}}^{\beta=6}. \quad (5.19)$$

Given $\sigma_{\dot{\nu}}$, we wish to estimate the spread of observed $\dot{\nu}$ values across different realisations of the noise process. We denote this spread by $\delta\dot{\nu}$. Over an observation of length T , $\delta\dot{\nu}$ can be estimated by calculating the dispersion in $\dot{\nu}$ due to the random walk, given by $\sigma_{\dot{\nu}} T^{1/2}$, and dividing this value by T , i.e.

$$\delta\dot{\nu} \sim \sqrt{\frac{f_{\text{yr}}^3 \nu^2 (2\pi)^6}{24\pi^2}} A_{\text{red}}^{\beta=6} T^{-1/2}. \quad (5.20)$$

If the $\dot{\nu}$ which we infer is to represent the *long-term* spin-down trend, $\delta\dot{\nu}$ ought to give an approximate lower bound on the reported uncertainty on $\dot{\nu}$ for pulsars with timing noise following equations (5.14)–(5.16). Any long-term secular $\dot{\nu}$ smaller than this value is liable to be washed out by the influence of the timing noise.

We compare the value of $\delta\dot{\nu}$ to $\Delta\dot{\nu}$, the estimated uncertainty on $\dot{\nu}$ returned by ENTERPRISE for all pulsars with $A_{\text{red}} > 10^{-11}$. As in Section 5.4.1, this restriction on A_{red} excludes pulsars with low-amplitude timing noise whose PSDs are poorly measured. For the TNF2 case, the median value of $\delta\dot{\nu}/\Delta\dot{\nu}$ is 2.79, while in the TNLONGF2 case it is 0.98. The TNF2 model tends to underestimate $\Delta\dot{\nu}$ relative to the expected spread in $\dot{\nu}$ due to timing noise, $\delta\dot{\nu}$, while the $\dot{\nu}$ uncertainty estimated by the TNLONGF2 model better reflects this variance.

We repeat this analysis for pulsars with $|\beta - 4| < 1$, suggesting a Langevin equation of the form

$$\frac{d\nu}{dt} = \dot{\nu}_0 + \xi(t), \quad (5.21)$$

with the strength of the $\xi(t)$ term now parametrised by σ_{ν} rather than $\sigma_{\dot{\nu}}$. By analogous arguments to the $|\beta - 6| < 1$ case, we find that the expected spread in observed $\dot{\nu}$ is given by (details are given in Appendix 5.A2)

$$\delta\dot{\nu} \sim 2 \sqrt{\frac{f_{\text{yr}} \nu^2 (2\pi)^4}{24\pi^2}} A_{\text{red}}^{\beta=4} T^{-3/2}. \quad (5.22)$$

Again comparing $\delta\dot{\nu}$ with $\Delta\dot{\nu}$, we find that the mean $\delta\dot{\nu}/\Delta\dot{\nu}$ in our sample is 0.49 and 0.40 for TNF2 and TNLONGF2 respectively — i.e., both timing noise models appear to overestimate the uncertainty in $\dot{\nu}$ relative to the variance due to timing noise, but the degree of the overestimation is less sensitive to f_{low} for $|\beta - 4| < 1$ than for $|\beta - 6| < 1$. Figure 5.5 shows histograms of $\log_{10}(\delta\dot{\nu}/\Delta\dot{\nu})$ for $|\beta - 6| < 1$ (top) and $|\beta - 4| < 1$ (bottom). Each panel shows two histograms of $\log_{10}(\delta\dot{\nu}/\Delta\dot{\nu})$, one for TNF2 and the other

for TNLONGF2. For $|\beta - 6| < 1$, there is a clear separation between the histograms for the two timing noise models, with the TN model tending to have larger values of $\delta\dot{\nu}/\Delta\dot{\nu}$, i.e. the spread due to timing noise exceeds the reported uncertainty, as noted above. For $|\beta - 4| < 1$, both models produce similar histograms and typically have $\delta\dot{\nu}/\Delta\dot{\nu} < 1$. These results broadly agree with the conclusions of Keith and Nițu (2023) – pulsars with $\beta \sim 6$ suffer more from mis-estimation of $\dot{\nu}$ uncertainty than those with $\beta \sim 4$, especially if a timing noise model with a low-frequency cutoff at $1/T_{\text{span}}$ is used.

5.5.2 | Measuring the braking index n_{pl}

The values of $\delta\dot{\nu}$ obtained in Section 5.5.1 can also be cast in terms of the measured braking index. We define

$$\delta n = \frac{\nu \delta\dot{\nu}}{\dot{\nu}^2}, \quad (5.23)$$

as the expected variation in measured n due to the influence of timing noise. This raises an interesting question: if for every pulsar in our dataset $1 \lesssim n_{\text{pl}} \lesssim 7$, for how many pulsars might we expect to measure n_{pl} , given our current dataset? And if timing noise currently confounds a measurement of n_{pl} , for how long do we need to observe to overcome this effect? Vargas and Melatos (2023a) derived a condition for when we measure $|n| \gg n_{\text{pl}}$ in terms of the strength of the noise process, the pulsar spin frequency and its first derivative, and the observing timespan; see equation (14) in Vargas and Melatos (2023a). Their condition follows from $\delta n > 1$. Here we write the equivalent condition on the noise strength for $|\beta - 6| < 1$ [which is a condition on $\sigma_{\dot{\nu}}$, see Section 5.5.1], viz.

$$\sigma_{\dot{\nu}} > 10^{-20} \left(\frac{\dot{\nu}}{10^{-12} \text{ Hz s}^{-1}} \right)^2 \left(\frac{\nu}{1 \text{ Hz}} \right)^{-1} \left(\frac{T_{\text{span}}}{10^8 \text{ s}} \right)^{1/2} \text{ s}^{-5/2}, \quad (5.24)$$

and $|\beta - 4| < 1$ [which is a condition on σ_{ν} , see Appendix 5.A2], viz.

$$\sigma_{\nu} > 10^{-12} \left(\frac{\dot{\nu}}{10^{-12} \text{ Hz s}^{-1}} \right)^2 \left(\frac{\nu}{1 \text{ Hz}} \right)^{-1} \left(\frac{T_{\text{span}}}{10^8 \text{ s}} \right)^{3/2} \text{ s}^{-3/2}. \quad (5.25)$$

These can be rephrased as conditions on $A_{\text{red}}^{\beta=6}$, viz.

$$A_{\text{red}}^{\beta=6} > 1.1 \times 10^{-10} \left(\frac{\dot{\nu}}{10^{-12} \text{ Hz s}^{-1}} \right)^2 \left(\frac{\nu}{1 \text{ Hz}} \right)^{-2} \left(\frac{T_{\text{span}}}{10^8 \text{ s}} \right)^{1/2}, \quad (5.26)$$

and $A_{\text{red}}^{\beta=4}$, viz.

$$A_{\text{red}}^{\beta=4} > 1.1 \times 10^{-9} \left(\frac{\dot{\nu}}{10^{-12} \text{ Hz s}^{-1}} \right)^2 \left(\frac{\nu}{1 \text{ Hz}} \right)^{-2} \left(\frac{T_{\text{span}}}{10^8 \text{ s}} \right)^{3/2}. \quad (5.27)$$

With equations (5.24)–(5.27) in hand, we ask: for which pulsars are the relevant inequalities *not* satisfied? That is, for which pulsars do we expect the variation in the measured value of n due to the influence of timing noise to be less than unity? Table 5.4 lists the two pulsars with $\delta n < 1$ over the timespan covered by the combined




Figure 5.5: Histograms comparing the expected variation in the measured value of $\ddot{\nu}$ due to timing noise, $\delta\ddot{\nu}$, against the uncertainty reported by ENTERPRISE, $\Delta\ddot{\nu}$. The results are divided into two cases depending on the estimated spectral index of the PSD, with $|\beta - 6| < 1$ in the top panel and $|\beta - 4| < 1$ in the bottom panel. Two histograms are shown in each panel, corresponding to the TN (blue) and TNLONG (orange) timing noise models.

Table 5.4: Pulsars with $\delta n < 1$, where δn is the expected variation due to timing noise in the measured value of n , as discussed in Section 5.5.1. The values of $\log_{10} A_{\text{red}}$, β , $\dot{\nu}$ and n are taken from the TNLONGF2 model, and the values of $\delta\dot{\nu}$ and δn are calculated according to equations (5.22) and (5.23). The quoted uncertainties on $\dot{\nu}$ and n are 95% credible intervals.

| PSR | $\log_{10} A_{\text{red}}$ | β | T_{span} (yr) | $\dot{\nu}$ (Hz s^{-2}) | $\delta\dot{\nu}$ (Hz s^{-2}) | n | δn |
|------------|----------------------------|---------|------------------------|------------------------------------|--|------------------------|------------|
| J0534+2200 | -9.22 | 4.08 | 1.9 | $(1.16 \pm 0.02) \times 10^{-20}$ | 3.6×10^{-23} | $2.54_{-0.03}^{+0.04}$ | 0.0079 |
| J1731-4744 | -9.78 | 3.66 | 8.2 | $(1.2 \pm 0.3) \times 10^{-24}$ | 2.9×10^{-26} | $26.7_{-6.6}^{+6.7}$ | 0.62 |




Figure 5.6: Distribution of T_{span} required to obtain $\delta n < 1$ [(equation 5.23)], for the cases where the spectral index of the timing noise PSD satisfies $|\beta - 6| < 1$ and $|\beta - 4| < 1$. The median required T_{span} is 2.2×10^7 yr for $|\beta - 6| < 1$, much longer than the 2.8×10^2 yr for $|\beta - 4| < 1$.

EW+NS dataset, using the parameters returned by the TNLONGF2 model. Note that we are not claiming to have measured a significant non-zero $\dot{\nu}$ in both pulsars — only PSR J0534+2200 favours an F2 model over all non-F2 models with a log Bayes factor greater than 5. The statement that $\delta n < 1$ is only a statement about the influence of the timing noise on the measurement of n , not whether a second frequency derivative is necessary to explain the observed data. The question of detection of non-zero $\dot{\nu}$ via model selection is discussed further in Section 5.5.3.

For all other pulsars we do not expect to be able to resolve $\dot{\nu}$ due to “canonical” secular spindown with the UTMOST-NS+EW data. However, δn is a decreasing function of T for both $|\beta - 4| < 1$ and $|\beta - 6| < 1$, so we expect that with a sufficiently long timing baseline, n_{pl} can be measured [again neglecting for simplicity any mean reversion (Vargas and Melatos, 2023a)]. Figure 5.6 shows histograms of the required T_{span} to obtain $\delta n < 1$ for the cases $|\beta - 6| < 1$ (blue) and $|\beta - 4| < 1$ (orange). Pulsars with $|\beta - 6| < 1$ require significantly longer T_{span} on average — the median required T_{span} is 3.9×10^7 yr, compared to 2.8×10^2 yr for $|\beta - 4| < 1$.

5.5.3 | Influence of frequency cut-off in timing noise models on the detection of $\dot{\nu}$

As discussed at the beginning of Section 5.5, if models with a $1/(2T_{\text{span}})$ cutoff are included in the set of possible timing models, PSR J0534+2200 is the only pulsar to show significant evidence for a $\dot{\nu} \neq 0$ (i.e. an F2 model is favoured over every non-F2 model with $\ln \mathcal{B} > 5$). However, had we not included the TNLONG models in our analyses, eight additional pulsars would have favoured an F2 model, as shown in Table 5.5. The values of n implied by these $\dot{\nu}$ measurements range between -2.4×10^4 and 5.0×10^4 . A comprehensive investigation into the reliability of model selection between $\dot{\nu} = 0$ and $\dot{\nu} \neq 0$ is

Table 5.5: Measurements of non-zero $\dot{\nu}$ when TNLONG models are excluded from consideration, as discussed in Section 5.5.3. The measured value of $\dot{\nu}$ is given in the second column, while the implied value of n [equation (5.13)] appears in the third. The estimated properties of the timing noise PSD, A_{red} and β , are given in the fourth and fifth columns. The model chosen according to the criteria defined in Section 5.3.2 is listed in the penultimate column, and the final column gives the log Bayes factor between this chosen model and the model with the highest evidence which does not contain a $\dot{\nu}$ term.

| PSR | $\dot{\nu}$ (Hz s ⁻²) | n | $\log_{10} A_{\text{red}}$ | β | Model | $\ln \mathcal{B}_{\dot{\nu}=0}^{\dot{\nu} \neq 0}$ |
|------------|---|------------------------------------|----------------------------|---------|-------|--|
| J0534+2200 | $1.16_{-0.01}^{+0.02} \times 10^{-20}$ | 2.54 ± 0.03 | -9.2 | 4.13 | TNF2 | 1.6×10^3 |
| J0835-4510 | $(1.23 \pm 0.01) \times 10^{-21}$ | 57.4 ± 0.6 | -8.46 | 7.45 | TNF2 | 350 |
| J1048-5832 | $(1.2 \pm 0.2) \times 10^{-22}$ | $25.1_{-4.2}^{+4.6}$ | -8.72 | 4.85 | TNF2 | 31.9 |
| J1522-5829 | $-4.14_{-0.05}^{+0.04} \times 10^{-25}$ | $(-6.0_{-0.7}^{+0.6}) \times 10^3$ | - | - | F2 | 7.38 |
| J1645-0317 | $(-1.3 \pm 1.3) \times 10^{-24}$ | $(-2.4_{-2.4}^{+2.4}) \times 10^4$ | -10.0 | 8.49 | TNF2 | 8.03 |
| J1709-4429 | $(2.1 \pm 0.2) \times 10^{-22}$ | $26.0_{-2.3}^{+2.7}$ | -9.61 | 6.43 | TNF2 | 40.4 |
| J1731-4744 | $(1.3 \pm 0.2) \times 10^{-24}$ | $27.5_{-4.8}^{+5.3}$ | -9.79 | 3.6 | TNF2 | 6.26 |
| J1740-3015 | $8.2_{-2.5}^{+2.6} \times 10^{-24}$ | $8.5_{-2.6}^{+2.7}$ | -9.01 | 5.27 | TNF2 | 7.39 |
| J1741-3927 | $1.0_{-0.1}^{+0.2} \times 10^{-24}$ | $(5.0 \pm 0.7) \times 10^4$ | -11.5 | 9.04 | TNF2 | 12.7 |

beyond the scope of this work. It is possible, of course, that excluding timing noise models with a cut-off at $1/(2T_{\text{span}})$ may lead to spurious measurements of $\dot{\nu}$, and including them may lead to spurious non-measurements, as the contribution to the residuals from the secular $\dot{\nu}$ is absorbed into the low-frequency component of the timing noise. For example, in the case of PSR J0534+2200, the log Bayes factor between TNF2 (the overall most favoured model) and TNLONG is 61, whereas the log Bayes factor between TNF2 and TN is 1.6×10^3 . This is a robust measurement of $\dot{\nu}$ either way, but F2 is favoured more, when TNLONG is excluded. However, the lower cut-off in the timing noise power spectrum is connected to the intrinsic physical processes at play (Melatos and Link, 2014), and is not known *a priori*. In particular there is no reason for it to be connected to the observing time span, and hence no reason to favour either TN or TNLONG. Consequently, although some of the measurements in Table 5.5 may be measurements of a secular $\dot{\nu}$ [whether due to the long-term spin down mechanism or other processes such as glitch recovery (Hobbs et al., 2010; Lower et al., 2021; Liu et al., 2024)] which are absorbed into the timing noise by the TNLONG model, it is impossible to know for sure, so we do not interpret these measurements further.

5.6 | Glitches

A significant motivation for the UTMOST-NS pulsar timing programme is to monitor the pulsars in the programme for glitching activity. With approximately 50 pulsars per day being observed, manually inspecting every new observation for a possible glitch and responding in a timely manner is not possible, and an automated solution to “online” glitch detection is implemented using the hidden Markov model (HMM)-based glitch detector⁴ introduced by Melatos et al. (2020). In addition, a subsequent “offline” glitch search combining data from UTMOST-EW and UTMOST-NS is carried out, again using the HMM glitch detector. After an introduction to the HMM glitch detector, these two glitch searches, and their results, are described in the rest of this section.

5.6.1 | Glitch detection with an HMM

We begin with a brief overview of how the HMM-based glitch detector operates. The reader is referred to Melatos et al. (2020) for a full description of the method, and Lower et al. (2021), Dunn et al. (2022b), and Dunn et al. (2023a) for previous examples of the use of this method to search successfully for glitches.

HMMs provide a convenient way to model systems in which the internal state of the system evolves in a Markovian (i.e. memoryless) fashion, and in which we are unable to observe the internal state directly, instead making proxy measurements which are connected probabilistically to the true internal states (i.e. the internal state is “hidden”) (Rabiner, 1989). In this paper, the internal state of the system (pulsar) is taken to be the tuple of the spin frequency and frequency derivative $q_i = (\nu_i, \dot{\nu}_i)$. The hidden state is indexed with the discrete variable i , as the space of possible hidden states is discretised on a bounded grid of ν and $\dot{\nu}$ values referred to as the “domain of interest” (DOI) (see Section 5.6.2). Time is discretised into a set of timesteps $\{t_1, t_2, \dots, t_{N_T}\}$, where each timestep

⁴github.com/ldunn/glitch_hmm

corresponds to a gap between consecutive ToAs. We assume that the hidden state of the pulsar is fixed during each gap, but may change at the boundary between gaps according to some prescribed Markovian model, discussed in more detail in Section 5.6.2, so that we have a sequence of hidden states $\{q(t_1), q(t_2), \dots, q(t_{N_T})\}$.

The connection between the observed ToAs and the $(\nu, \dot{\nu})$ states is provided by the emission likelihood which is the probability of observing a particular gap between consecutive ToAs, given the pulsar occupies some hidden state q_i during this gap. Explicitly, the emission likelihood takes the form of a von Mises distribution,

$$\Pr(z | q_i) = \frac{\exp[\kappa \cos(2\pi\Phi)]}{2\pi I_0(\kappa)}, \quad (5.28)$$

where Φ denotes the number of turns accumulated during a gap of length z given the pulsar is in the hidden state q_i , and κ is a parameter known as the ‘‘concentration’’, which parametrises the uncertainty due to both ToA measurement uncertainty and the discretisation of the space of hidden states. The reader is referred to Section 3 of Dunn et al. (2022b) for further discussion of the forms of Φ and κ . The emission likelihood peaks, where the fractional part of Φ is zero, i.e. choices of ν and $\dot{\nu}$ for which an integer number of turns elapse during the ToA gap.

In this formalism it is straightforward to compute the model evidence $\Pr(D | \mathcal{M})$, where D denotes the observed data and \mathcal{M} is a chosen model, which among other things incorporates assumptions about the Markovian transitions from one time step to the next. This leads to a method of glitch detection based on Bayesian model selection, wherein we compare a model in which the transition dynamics include only an assumed Markov process representing the timing noise, denoted by \mathcal{M}_0 , against a series of models which include both the Markov process and the possibility of a discrete jump in frequency and frequency derivative (i.e. a glitch) during the k th ToA gap, denoted by $\mathcal{M}_1(k)$. We compute the Bayes factors

$$K_1(k) = \frac{\Pr[D | \mathcal{M}_1(k)]}{\Pr[D | \mathcal{M}_0]} \quad (5.29)$$

for all k . If the maximum value of $K_1(k)$, denoted by $K_1^* = K_1(k^*)$, exceeds a pre-determined threshold, here taken to be $10^{1/2}$, we flag a glitch candidate. Multiple glitches in one dataset can be handled via a greedy procedure; see Section 4.2 of Melatos et al. (2020) for details.

We measure the properties of a detected glitch using the forward-backward algorithm (Rabiner, 1989) to efficiently compute the posterior distribution, i.e. $\gamma_{q_i}(t_n) = \Pr[q(t_n) = q_i | D]$. This distribution can be marginalised over either $\dot{\nu}$ or ν to obtain the posterior distribution on ν or $\dot{\nu}$ alone, respectively. Hence we can compute trajectories in ν or $\dot{\nu}$ by taking the sequence of a *a posteriori* most probable hidden states, and from these trajectories we may compute e.g. the size of the frequency jump associated with a glitch.

As an illustrative example, Figure 5.7 shows the results of the HMM analysis of a new small glitch in PSR J1902+0615 (see Section 5.6.5). The top panel shows the log Bayes factors obtained as part of the model selection procedure. The blue points are the values of $K_1(k)$ [see equation (5.29)], which peak at $k^* = 57$ (i.e. the 57th ToA gap) with a maximum value of $\ln K_1^* = 15.9$. The orange points show the results of the second




Figure 5.7: Illustrative example of a glitch analysis for PSR J1902+0615 using the HMM pipeline (Melatos et al., 2020), discussed in Section 5.6.1. The left panel shows the results of the model selection procedure, with log Bayes factors $K_1(k)$ (comparing one-glitch models against the no-glitch model; k indexes the ToA gap containing a glitch) and $K_2(57, k)$ (comparing two-glitch models against the best one-glitch model) plotted versus k (matched to MJD) in blue and orange respectively. The peak one-glitch log Bayes factor is 15.9, well above the threshold of 1.15, while the peak two-glitch log Bayes factor is 0.32. The middle and right panels show heatmaps of the posterior distribution of the hidden states ν and $\dot{\nu}$, with the maximum a posteriori tracks overlaid in red. The estimated size of the glitch based on the ν posterior distribution is $(1.1 \pm 0.1) \times 10^{-9}$ Hz. There is no discernible jump in $\dot{\nu}$. The vertical dashed lines indicate the location of the glitch. The values of γ_ν and $\gamma_{\dot{\nu}}$ are normalised so they sum to unity at each timestep.

iteration of the greedy algorithm, where we compare the two-glitch model $\mathcal{M}_2(k^*, k)$ against the one-glitch model $\mathcal{M}_1(k^*)$ for all $k \neq k^*$. The largest log Bayes factor $K_2(k^*, k)$ obtained in this second step is 0.32, below the threshold $\ln 10^{1/2} = 1.15$. Thus the model selection terminates and we take $\mathcal{M}_1(57)$ as our preferred model. The middle and bottom panels of Figure 5.7 show heatmaps of the posterior distributions of the hidden states ν and $\dot{\nu}$. The maximum *a posteriori* values are overlaid as red tracks, and the location of the glitch is indicated by the vertical dashed line. The jump in ν is clear, occurring around MJD 58590 with size $(1.1 \pm 0.1) \times 10^{-9}$ Hz. No change in $\dot{\nu}$ is apparent at the time of the glitch.

5.6.2 | HMM parameters

Completely specifying the HMM requires choosing a number of parameters controlling what set of $(\nu, \dot{\nu})$ pairs are allowed (i.e. the DOI), and how the hidden states evolve over time due to timing noise. In this section we describe some general considerations influencing both choices. The specifics of how the parameters are set for the online and offline glitch searches are discussed in Section 5.6.3 and Appendix 5.A3 respectively.

Domain of interest

The DOI is the set of hidden states available to the HMM in a given analysis. The hidden states are tuples of frequency and frequency derivative $(\nu, \dot{\nu})$. For practical reasons we define the hidden states *nu* ($\dot{\nu}$) to be the difference between the spin frequency (frequency derivative) and its approximate value as determined by an initial least-squares fit, denoted by ν_{LS} ($\dot{\nu}_{\text{LS}}$). The least-squares values are defined at a reference epoch T_0 . The phase accumulated during a gap of length z ending at time t_n is thus given by

$$\Phi(z, t_n) = [\nu_{\text{LS}} + \nu + \dot{\nu}_{\text{LS}}(t_n - T_0)]z - \frac{1}{2}(\dot{\nu}_{\text{LS}} + \dot{\nu})z^2, \quad (5.30)$$

where the second term includes a minus sign because we are performing a backwards Taylor expansion (see Section 3.3 of Melatos et al. 2020). A secular second frequency derivative $\ddot{\nu}_{\text{LS}}$ can also be included, if necessary. Note that because the spin-down behaviour is included via the secular terms, we typically take the DOI to be symmetric about $(\nu, \dot{\nu}) = (0, 0)$.

In all cases in this work the DOI is laid out as a uniform grid in $(\nu, \dot{\nu})$. Specifying the DOI thus amounts to specifying the bounds and grid spacing in each parameter. This varies on a case-by-case basis and is discussed in more detail in Section 5.6.3 for the online detection case, and Appendix 5.A3 for the offline detection case. Here we briefly introduce the notation used to describe the DOI for ν ; for $\dot{\nu}$ the notation is identical but replacing ν with $\dot{\nu}$. The bounds of the DOI in ν are denoted by $[\nu_-, \nu_+]$. As mentioned, the DOI is typically symmetric about $(0, 0)$, whereupon we have $\nu_- = -\nu_+$ and write the bounds as $[-\nu_{\pm}, \nu_{\pm}]$. The grid spacing is denoted by ϵ_ν , i.e. ν takes on values $\{\nu_-, \nu_- + \epsilon_\nu, \nu_- + 2\epsilon_\nu, \dots, \nu_+\}$.

Timing noise models

In principle the HMM specified in Section 5.6.1 can be used with various models for the underlying hidden state evolution, which amounts to specifying the model for the timing

noise in the pulsar. This matches the situation in traditional, non-HMM glitch searches. Many past analyses assume that the stochastic evolution is driven by a white noise term in the second frequency derivative, i.e. a random walk in $\dot{\nu}$. This model is identical to the one specified by equations (5.14)–(5.16). We call it RWF1. The form of the transition matrix $A_{q_i q_j} = \Pr[q(t_{n+1}) = q_j \mid q(t_n) = q_i, o(t_n) = z]$ is a multivariate Gaussian with the covariances

$$\text{cov}(\nu, \nu) = \sigma_{\dot{\nu}}^2 z^3 / 3, \quad (5.31)$$

$$\text{cov}(\dot{\nu}, \nu) = \sigma_{\dot{\nu}}^2 z^2 / 2, \quad (5.32)$$

$$\text{cov}(\dot{\nu}, \dot{\nu}) = \sigma_{\dot{\nu}}^2 z. \quad (5.33)$$

The reader is referred to Section 3.4 of Melatos et al. (2020) for the full expression for $A_{q_i q_j}$.

As discussed in Section 5.5.1, the RWF1 model gives rise to phase residuals with a PSD which has a spectral index of -6 . While many pulsars do exhibit residuals of this sort, a significant fraction also exhibit shallower PSDs. This motivates us to introduce a second model for timing noise in the HMM, this time driven by a white noise term in the frequency derivative, which corresponds to a random walk in ν and which gives rise to residuals with a spectral index of -4 . We call this model RWF0, and its Langevin equation is

$$\frac{d\nu}{dt} = \dot{\nu}_0 + \xi(t), \quad (5.34)$$

where $\xi(t)$ again satisfies equations (5.15) and (5.16), except that the parameter determining the noise amplitude is labelled σ_{ν} , instead of $\sigma_{\dot{\nu}}$. Writing down the covariances of the Gaussian describing the transition matrix in this case requires some care — the autocorrelation of $\dot{\nu}$ diverges, with $\langle \dot{\nu}(t) \dot{\nu}(t') \rangle \propto \delta(t - t')$. To make progress, we note that our measurements (the ToAs) tell us only about the average behaviour of the pulsar during a ToA gap, not the value of $d\nu/dt$ at any given instant (e.g. the end of a ToA gap). The delta-correlated nature of $\dot{\nu}(t)$ means that its value at any given instant is not correlated with its value at nearby times, and hence our measurements, which necessarily take place over a finite timespan and involve an averaging process, contain no information about $\dot{\nu}(t)$ at any particular instant. The hidden state variable $\dot{\nu}$ is therefore more usefully identified as the effective rate of change in ν over a ToA gap, i.e. $z^{-1} [\nu(t_2) - \nu(t_1)]$. The dispersion in the hidden state variable $\dot{\nu}$ is then better characterised by the dispersion in ν due to the random walk, which is $\sigma_{\nu} z^{1/2}$, divided by z to obtain a dispersion in the average $\dot{\nu}$ over the ToA gap⁵ — hence $\text{cov}(\dot{\nu}, \dot{\nu}) = \sigma_{\nu}^2 / z$. Under this interpretation, $\dot{\nu}$ is a linear function of $\nu(t_1)$ and $\nu(t_2)$. This presents an issue for the HMM glitch detector as currently formulated, as it implies a singular covariance matrix between ν and $\dot{\nu}$. The transition matrix $A_{q_i q_j}(z)$ involves the inverse of this covariance matrix and is therefore undefined. To circumvent this issue, as an approximation we neglect the correlation between ν and $\dot{\nu}$. The covariances for the RWF0 model are therefore taken to be

$$\text{cov}(\nu, \nu) = \sigma_{\nu}^2 z, \quad (5.35)$$

$$\text{cov}(\dot{\nu}, \nu) = 0, \quad (5.36)$$

$$\text{cov}(\dot{\nu}, \dot{\nu}) = \sigma_{\nu}^2 / z. \quad (5.37)$$

⁵This is similar to the argument in Section 5.5.1 to derive equation (5.20).

Table 5.6: HMM glitch detector parameters used in the online glitch detection analyses described in Section 5.6.3. Most parameters are fixed, but σ depends on the average ToA gap length in a 30-ToA window, denoted by \bar{z} .

| Parameter | Value |
|------------------------|---|
| ν_{\pm} | 1×10^{-6} Hz |
| ϵ_{ν} | 1×10^{-8} Hz |
| $\dot{\nu}_{\pm}$ | 2×10^{-14} Hz s ⁻¹ |
| $\epsilon_{\dot{\nu}}$ | 5×10^{-16} Hz s ⁻¹ |
| Timing noise model | RWF1 |
| $\sigma_{\dot{\nu}}$ | $\max(\epsilon_{\dot{\nu}} \bar{z}^{-1/2}, 10^{-21} \text{ Hz s}^{-3/2})$ |

An alternative solution is to reformulate the HMM to track ν alone, which we defer to future work.

5.6.3 | Online glitch detection

For every new observation which produced a good ToA, an HMM analysis of the most recent 30 ToAs⁶ for that pulsar is automatically run with low latency (~ 1 hour) to check for a potential glitch. In the event of a possible detection (any glitch candidate with a log Bayes factor greater than $10^{1/2}$; see Section 5.6.1), an email alert is generated and the data are further scrutinised to determine whether there is compelling evidence that a glitch has occurred, whereupon a public alert is sent out manually. In total, 2963 online HMM analyses produce a glitch candidate. In approximately 80% of cases the glitch candidate can be ascribed to RFI contamination which is not completely removed by the timing pipeline (see Section 5.2.3) and leads to a corrupted ToA. For 81% of the remaining 539 candidates, there is no obvious contamination in the data, but further investigation – either visual inspection of the timing residuals or re-analysis using the HMM with a finer grid in ν and $\dot{\nu}$ – reveals no compelling glitch candidate, and hence no public alert is sent out. In three pulsars, clear glitches occurred while the online glitch detection pipeline was operational (approximately October 2021 through June 2023), namely PSRs J0835–4510, J0742–2822, and J1740–3015. In all three objects the glitches were large, and easily detected by the online pipeline (Dunn et al., 2021b; Dunn et al., 2022a; Dunn et al., 2023b) – the remainder of the glitch candidates generated by the online pipeline contain one of these three glitches within their analysis window.

For simplicity, the online glitch searches use a uniform set of parameters for every pulsar, with values listed in Table 5.6. The timing noise model is fixed to be RWF1. The one parameter which is allowed to vary is $\sigma_{\dot{\nu}}$, which depends on the average ToA gap length within the 30-ToA window, denoted by \bar{z} .

The procedure adopted in the offline glitch searches described in Section 5.6.4 and Appendix 5.A3 was not in hand at the time that the online pipeline was in operation. This procedure allows the offline search to achieve better sensitivity to small glitches. Estimates of the sensitivity of both the online and offline pipelines are discussed in Section 5.6.6.

⁶The window length is chosen arbitrarily, but synthetic data tests show that glitch detection probability is a weak function of the number of ToAs; see Appendix G of Melatos et al. (2020).

5.6.4 | Offline glitch detection

The offline glitch searches are conducted on the same combined EW-NS datasets analysed in Section 5.3. In contrast to earlier HMM analyses involving a large number of objects (Lower et al., 2021; Dunn et al., 2022b), here we make explicit use of the timing noise information returned by the ENTERPRISE analyses of Section 5.3 to better match the timing noise model included in the HMM with the observed data. The timing noise information returned by ENTERPRISE consists of the model selection results, and the values of A_{red} and β in (5.3), if applicable. The two models available are RWF0 and RWF1, discussed in Section 5.6.2. The models generate phase residual PSDs with spectral indices of -4 and -6 respectively. Based on the value of β returned by ENTERPRISE, we assign a timing noise model to each pulsar, and based on the value of A_{red} we set the amplitude of the white noise driving term. The detailed procedure for choosing HMM parameters is described in Appendix 5.A3.

A number of glitches are already identified prior to running the offline glitch search, by the online search (see Section 5.6.3), by eye during the initial timing process using TEMPO2 and ENTERPRISE, or by other authors, e.g. Lower et al. (2020), Zubieta et al. (2024), and Keith et al. (2024). When the glitch is large ($\Delta\nu \gtrsim 10^{-7}$ Hz), the datasets passed to the HMM pipeline are split into pre- and post-glitch sections to avoid the need for a large ν range to accommodate the glitch. In one case, PSR J1048–5832, we divide the full dataset in two equal halves not because of the presence of a glitch but because the strong timing noise makes an analysis of the full dataset prohibitively expensive. In this case the two halves overlap by 60 days in order to ensure that a glitch close to the boundary between the two subsets is not missed. Datasets containing smaller glitches are analysed in full, in order to check with the HMM that these by-eye identifications are not timing noise fluctuations which have been miscategorised.

The complete list of glitch candidates identified by the offline HMM pipeline is given in Table 5.7. When multiple glitch candidates in a single pulsar are returned, we list only the most significant candidate. These candidates are subjected to a simple veto procedure, whereby the dataset containing the candidate is re-analysed with the two ToAs bracketing the glitch candidate removed (Dunn et al., 2022b). The veto is designed to exclude the case where a single ToA is displaced from the overall trend for non-astrophysical reasons (e.g. RFI contamination, clock errors), thereby generating a glitch candidate in the ToA gap immediately before or after the displaced ToA. If the re-analysis fails to return a candidate above the detection threshold, the initial candidate is vetoed. After vetos, the offline pipeline does not produce any additional glitch candidates over those identified by the online pipeline and by-eye inspection of the timing residuals. However it does produce tighter upper limits on the size of undetected glitches than the online pipeline, as discussed in Section 5.6.6.

5.6.5 | Glitches in the sample

In addition to those glitches detected during the UTMOST-NS timing programme and those already known from previous analyses, several glitches are identified when inspecting the timing residuals by eye to prepare for the ENTERPRISE analysis of the combined datasets.

The full list of the 17 glitches known in the combined EW+NS dataset is given in Table

Table 5.7: Glitch candidates identified by the offline HMM glitch search described in Section 5.6.4. The detection threshold is $\ln K_{\text{th}} = 1.15$. Note that large ($\Delta\nu \gtrsim 10^{-7}$ Hz) glitches are not included here, as the datasets for those pulsars are broken into pre- and post-glitch sections to allow for smaller DOIs and hence more sensitive searches. We indicate whether the glitch was previously identified, either as part of the UTMOST-NS online glitch detection pipeline or through by-eye inspection of the timing residuals, and whether the glitch candidate is vetoed by the procedure described in Section 5.6.4.

| PSR | Glitch epoch (MJD) | $\ln K_1^*$ | Previously identified? | Vetoed? |
|------------|--------------------|-------------------|------------------------|---------|
| J0835–4510 | 57756–57763 | 1.3 | N | Y |
| J0908–4913 | 60092–60119 | 8.1 | N | Y |
| J1136–5525 | 60078–60115 | 4.4 | N | Y |
| J1257–1027 | 58666–58697 | 1.3×10^2 | Y | N |
| J1326–6700 | 57438–57520 | 4 | N | Y |
| J1453–6413 | 59327–59359 | 4.4×10^2 | Y | N |
| J1703–4851 | 58564–58668 | 2.7×10^2 | Y | N |
| J1836–1008 | 59427–59458 | 2.1×10^2 | Y | N |
| J1902+0615 | 58624–58655 | 7.2 | Y | N |
| J1909–3744 | 57267–57325 | 1.7 | N | Y |

5.8. The glitch in PSR J1902+0615 is the only one which has not previously been identified in the literature. The glitch parameter estimates and their 95% credible intervals are derived from the ENTERPRISE analyses. To determine whether there is support for a decaying component, we perform the model selection procedure outlined in Section 5.3.2 twice for each glitch: once with a $\Delta\nu_d$ prior which is log-uniform on $[10^{-10} \text{ Hz}, 10^{-4} \text{ Hz}]$ and again with $\Delta\nu_d$ fixed at zero. If the log Bayes factor between the favoured model with a floating $\Delta\nu_d$ and the $\Delta\nu_d = 0$ model exceeds 5, we adopt the model with a decaying component, and report constraints on the decay parameters. Previous analyses have reported evidence for multiple exponential components with different timescales (e.g. McCulloch et al. 1987; Wong et al. 2001; Dodson et al. 2002; Shannon et al. 2016; Liu et al. 2024). A comprehensive search for additional decaying components is beyond the scope of this work.

In the remainder of this section we discuss each of the detected glitches in detail.

PSR J0742–2822

A new glitch was reported in PSR J0742–2822 by Shaw et al. (2022) in September 2022, and soon after confirmed by UTMOST-NS (Dunn et al., 2022a) as well as Grover et al. (2022) and Zubieta et al. (2022b). It is confidently detected by the online pipeline in the first post-glitch observation at UTMOST-NS, taken on September 22 (MJD 59844.9), with $\ln K_1^* = 123$. With a permanent frequency step of $\Delta\nu_p/\nu = (4284 \pm 2) \times 10^{-9}$, it is by far the largest glitch out of the nine known in this pulsar; the previous largest glitch had an amplitude of $\Delta\nu/\nu = 92(2) \times 10^{-9}$ (Espinoza et al., 2011). We detect a decaying component ($\ln \mathcal{B}_{\Delta\nu_d=0}^{\Delta\nu_d \neq 0} = 104$) with $\Delta\nu_d/\nu = (16.2 \pm 1.2) \times 10^{-9}$ and $\tau_d = 14_{-2}^{+3}$ d. Zubieta et al. (2024) also report evidence for a decaying component, although with a somewhat longer timescale, viz. $\Delta\nu_d/\nu = (19.16 \pm 0.04) \times 10^{-9}$ and $\tau_d = 33.4 \pm 0.5$ d.

Table 5.8: Timing properties of all glitches analysed in this work. Parameter estimates are from ENTERPRISE analyses, uncertainties are 95% credible intervals. The glitches are discussed in detail in Section 5.6.5.

| PSR | Glitch epoch (MJD) | $\Delta\nu_r/\nu \times 10^9$ | $\Delta\dot{\nu}_p/\dot{\nu} \times 10^3$ | $\Delta\nu_d/\nu \times 10^9$ | τ_d (d) |
|------------|--------------------|-------------------------------|---|-------------------------------|---------------------|
| J0742–2822 | 59840.0 | 4284 ± 2 | 10 ± 4 | $16.2^{+1.3}_{-1.2}$ | 14^{+3}_{-2} |
| J0835–4510 | 57734.4855 | 1432.4 ± 0.4 | $6.9^{+0.4}_{-0.5}$ | $9.2^{+1.7}_{-1.4}$ | $1.9^{+0.7}_{-0.6}$ |
| J0835–4510 | 58515.0 | 2479^{+2}_{-2} | 7.0 ± 0.8 | 18 ± 2 | $7.8^{+2.2}_{-2.3}$ |
| J0835–4510 | 59417.2108 | 1236.9 ± 0.3 | 5.8 ± 0.3 | $6.5^{+1.4}_{-1.3}$ | $2.9^{+0.8}_{-0.7}$ |
| J0908–4913 | 58765.06 | 22.0 ± 0.2 | 0.4 ± 0.8 | – | – |
| J1105–6107 | 58581.0 | 1171.4 ± 0.5 | $2.2^{+1.8}_{-1.9}$ | – | – |
| J1257–1027 | 58649.3 | 3.07 ± 0.01 | 2.3 ± 0.4 | – | – |
| J1453–6413 | 59326.0 | 0.9 ± 0.1 | $0.0^{+0.2}_{-0.1}$ | – | – |
| J1703–4851 | 58522.0 | 11.3 ± 0.1 | 3.1 ± 0.3 | – | – |
| J1709–4429 | 58178.0 | 2415 ± 6 | 4.3 ± 0.6 | 20.0 ± 5 | 55^{+16}_{-17} |
| J1731–4744 | 57984.260363 | 3133^{+10}_{-9} | $0.4^{+0.2}_{-0.3}$ | $16.6^{+9.1}_{-9.3}$ | 613^{+380}_{-353} |
| J1740–3015 | 57468.0 | 229 ± 1 | 0.7 ± 0.5 | – | – |
| J1740–3015 | 58241.0 | 837^{+2}_{-3} | 1.3 ± 0.5 | $5.3^{+3.8}_{-5.3}$ | 56^{+203}_{-56} |
| J1740–3015 | 59945.0 | 328 ± 1 | $1.2^{+0.7}_{-0.6}$ | – | – |
| J1803–2137 | 58958.0 | 4592^{+99}_{-93} | $5.2^{+3.0}_{-3.5}$ | – | – |
| J1836–1008 | 59400.0 | 36 ± 3 | $1.4^{+4.3}_{-5.3}$ | – | – |
| J1902+0615 | 58590.0 | 0.76 ± 0.02 | 0.1 ± 0.1 | – | – |

The most recent large glitch in PSR J0835–4510 is the first one to be detected during the UTMOST-NS observing run, occurring in July 2021. It was first reported by Sosa-Fiscella et al. (2021), and subsequently confirmed by a number of other observatories (Dunn et al., 2021b; Olney, 2021; Singha et al., 2021b). Its characteristics are typical for Vela glitches, with a large permanent frequency jump of $\Delta\nu_p/\nu = (1236.9 \pm 0.3) \times 10^{-9}$. We find strong evidence for a decaying component ($\ln \mathcal{B}_{\Delta\nu_d=0}^{\Delta\nu_d \neq 0} = 109$), with $\Delta\nu_d/\nu = (6.5_{-1.3}^{+1.4}) \times 10^{-9}$ and $\tau_d = 2.9_{-0.7}^{+0.8}$ d. Zubieta et al. (2024) reported three distinct recovery components, with timescales of 0.994(8) d, 6.400(2) d, and 535(8) d. The reported fractional amplitudes of the two shortest-timescale recoveries are $9(1) \times 10^{-9}$ and $3(1) \times 10^{-9}$ respectively, while the long-timescale recovery has a larger reported amplitude of 512×10^{-9} . Our detected decaying component aligns roughly with the two short-timescale components, but does not fall within the error bars of either of them.

Apart from the July 2021 glitch, two earlier Vela glitches are also covered by our dataset and have been discussed by a number of previous authors (Sarkissian et al., 2017; Palfreyman et al., 2018; Sarkissian et al., 2019; Kerr, 2019; Lopez Armengol et al., 2019; Lower et al., 2020; Gügercinoğlu et al., 2022). In the case of the glitch at MJD 57734 we detect a significant ($\ln \mathcal{B}_{\Delta\nu_d=0}^{\Delta\nu_d \neq 0} = 221$) decaying component with a large amplitude of $\Delta\nu_d/\nu = 9.2_{-1.4}^{+1.7} \times 10^{-9}$ and a short timescale of $\tau_d = 1.9_{-0.6}^{+0.7}$ d. This is approximately consistent with the analysis of Sarkissian et al. (2017), who found $\Delta\nu_d/\nu = (11.5 \pm 0.7) \times 10^{-9}$ and $\tau_d = (0.96 \pm 0.17)$ d. We also find strong evidence for a recovery in the glitch at MJD 58515 ($\ln \mathcal{B}_{\Delta\nu_d=0}^{\Delta\nu_d \neq 0} = 88$), with $\Delta\nu_d/\nu = 18 \pm 2 \times 10^{-9}$ and $\tau_d = 7.8_{-2.3}^{+2.2}$ d.

Previous Vela glitches have been observed to recover on a wide range of time scales, from minutes to hundreds of days (McCulloch et al., 1990; Shannon et al., 2016; Ashton et al., 2019a), with several previous studies reporting an intermediate recovery timescale on the order of days (Flanagan, 1990; McCulloch et al., 1990; Dodson et al., 2002; Sarkissian et al., 2017). In all three glitches analysed here we find strong evidence for a recovering component on this intermediate timescale and with amplitudes $\sim 10^{-7}$ Hz, similar to previous works. The recovery fractions $Q = \Delta\nu_d/(\Delta\nu_p + \Delta\nu_d)$ for the three glitches in our dataset are (from oldest to newest) 6.5×10^{-3} , 7.1×10^{-3} , and 5.3×10^{-3} , compared to 6.5×10^{-3} as reported by McCulloch et al. (1990) and 5.5×10^{-3} as reported by Dodson et al. (2002). An intermediate-timescale decaying component with a recovery fraction of approximately 5×10^{-3} appears to be a common feature of large Vela glitches.

This glitch was previously detected and discussed by Lower et al. (2019). It is the only glitch known in PSR J0908–4913. Our recovered amplitude of $\Delta\nu_p/\nu = (22.0 \pm 0.2) \times 10^{-9}$ is only marginally different from the value inferred by Lower et al. (2019) of $\Delta\nu_p/\nu = (21.7 \pm 0.1) \times 10^{-9}$. Like Lower et al. (2019) we are unable to distinguish $\Delta\dot{\nu}_p$ from zero within the measurement uncertainty. Lower et al. (2019) did not report an exponential recovery, and we find no evidence for a decaying component.

PSR J1105–6107

This glitch occurred in April 2019, between MJD 58580 and MJD 58584. It occurred while UTMOST-EW observations were still ongoing – the last UTMOST-EW observation of this pulsar was taken on April 28 2020. It was not reported in previous analyses (Lower et al., 2020; Dunn et al., 2022b) because it came after the cut-off for the first data release. However, a measurement was reported by Abbott et al. (2022a) based on UTMOST-EW data, as part of a search for continuous gravitational wave emission in the aftermath of the glitch. Although there are no UTMOST-NS ToAs available for this pulsar, here we re-analyse the UTMOST-EW data covering this glitch to provide updated parameter and uncertainty estimates. We find $\Delta\nu_p/\nu = (1171.4 \pm 0.5) \times 10^{-9}$ and $\Delta\dot{\nu}_p/\dot{\nu} = 2.2_{-1.9}^{+1.8} \times 10^{-3}$. This is the largest glitch reported in PSR J1105–6107; the next largest has a magnitude of $\Delta\nu_p/\nu = (971.7 \pm 0.5) \times 10^{-9}$ (Yu et al., 2013). As in Abbott et al. (2022a), we do not find evidence for a decaying component.

PSR J1257–1027

This glitch was reported by Lower et al. (2020) and is the only known glitch in this pulsar. Our recovered amplitude of $\Delta\nu_p/\nu = (3.07 \pm 0.01) \times 10^{-9}$ is consistent with the Lower et al. (2020) value of $\Delta\nu_p/\nu = 3.20_{-0.57}^{+0.16} \times 10^{-9}$. Lower et al. (2020) reported an upper limit on the change in frequency derivative, $\Delta\dot{\nu}_p/\dot{\nu} \lesssim 286 \times 10^{-3}$. With our extended dataset, we measure $\Delta\dot{\nu}_p/\dot{\nu} = 2.3 \pm 0.4$. Lower et al. (2020) did not detect an exponential recovery, and we find no evidence for a decaying component.

PSR J1453–6413

This glitch is identified by eye as occurring during the gap between the UTMOST-EW and UTMOST-NS datasets, between MJDs 58927 and 59326. It was subsequently confirmed in the offline glitch search, being confidently detected with $\ln K_1^* = 2.6 \times 10^2$. It was reported by Li et al. (2023) as occurring at MJD 59060(12) with $\Delta\nu/\nu = 1.180(7) \times 10^{-9}$, and by Keith et al. (2024) as occurring at MJD 59015 with $\Delta\nu/\nu = (1.14 \pm 0.13) \times 10^{-9}$. Our value of $\Delta\nu_p/\nu = (0.9 \pm 0.1) \times 10^{-9}$ is in mild tension with these values, and our limit on $\Delta\dot{\nu}_p/\dot{\nu} \lesssim 0.2 \times 10^{-3}$ is also in tension with the value measured by Li et al. (2023) of $0.50(3) \times 10^{-3}$; Keith et al. (2024) did not report a measurement of $\Delta\dot{\nu}$. It is likely that a combination of the long observing gap in the UTMOST data and differences in timing noise modelling account for these tensions. Li et al. (2023) employ a method based on a Cholesky decomposition (Coles et al., 2011; Dang et al., 2020) to model the red noise in the dataset, rather than the Fourier decomposition used in this work. We remind the reader that the results of any glitch search are conditional on the phase model and indeed the user-selected false alarm and false dismissal probabilities.

PSR J1703–4851

This glitch was first reported by Lower et al. (2020) and is the only known glitch in this pulsar. Our recovered glitch amplitude is somewhat smaller than the value reported by Lower et al. (2020), viz. $\Delta\nu_p/\nu = (11.3 \pm 0.1) \times 10^{-9}$ rather than $\Delta\nu_p/\nu = 19.0_{-0.7}^{+1.0} \times 10^{-9}$. We also find a significantly different value of $\Delta\dot{\nu}/\dot{\nu} = (3.1 \pm 0.3) \times 10^{-3}$ as compared to

the Lower et al. (2020) value of $292_{-53}^{+38} \times 10^{-3}$. The updated measurement of $\Delta\dot{\nu}$ is largely enabled by the addition of the NS data – the EW data used in Lower et al. (2020) are quite sparse post-glitch, with only four post-glitch ToAs available, covering six months. Future timing is likely to further refine this measurement, as the timing noise in this pulsar is low (a NLL model is preferred for our dataset), so measurements of the glitch parameters are not confounded by timing noise as the timing baseline is lengthened.

PSR J1709–4429

The glitch at MJD 58178 was first reported by Lower et al. (2018), with $\Delta\nu/\nu = (52.4 \pm 0.1) \times 10^{-9}$. This glitch amplitude was later discovered to be in error in the course of a re-analysis using the HMM glitch detector, and revised to $\Delta\nu/\nu = (2432.2 \pm 0.1) \times 10^{-9}$ by Dunn et al. (2021a). Although Dunn et al. (2021a) found no evidence for a decaying component, here we find strong evidence ($\ln \mathcal{B}_{\nu_d=0}^{\nu_d \neq 0} = 39$) for a decaying term with $\Delta\nu_d/\nu = (20 \pm 5) \times 10^{-9}$ and $\tau_d = 55 \pm 16$ d. The discrepancy may be due to the difference in the available post-glitch data – Dunn et al. (2021a) used data from the first UTMOST data release (Lower et al., 2020) which extends up to MJD 58418, 240 days after the glitch, while this analysis includes 710 days of post-glitch data before a long observing gap in early 2020. The total glitch amplitude $(\Delta\nu_p + \Delta\nu_d)/\nu = (2434 \pm 10) \times 10^{-9}$ is consistent with the value reported by Dunn et al. (2021a). The fractional change in frequency derivative reported by Dunn et al. (2021a) of $\Delta\dot{\nu}_p/\dot{\nu} = (4.7 \pm 0.3) \times 10^{-3}$ is consistent with the value obtained in this analysis, viz. $(4.3 \pm 0.6) \times 10^{-3}$.

PSR J1731–4744

The glitch at MJD 57984 was first reported by Jankowski et al. (2017), and is the largest glitch ever observed in this pulsar. None of the previous works analysing this glitch (Jankowski et al., 2017; Lower et al., 2020; Basu et al., 2020) have detected a post-glitch recovery. In our analysis we find moderate evidence ($\ln \mathcal{B}_{\nu_d=0}^{\nu_d \neq 0} = 7.6$) for the presence of a decaying component with $\Delta\nu_d/\nu = 16.6_{-9.3}^{+9.1} \times 10^{-9}$ and $\tau_d = 613_{-353}^{+380}$ d. It is possible that our extended dataset relative to other studies allows for better characterisation of the recovery, given the long timescale inferred. However, given that there is significant timing noise in this pulsar ($\log_{10} A_{\text{red}} = -9.7$, $\beta = 3.7$) it is also possible that there is confusion between the timing noise and a putative long-timescale recovery process, which would explain the large fractional error bar on the timescale.

PSR J1740–3015

PSR J1740–3015 most recently suffered a glitch in December 2022, first reported by Zubieta et al. (2022a) and subsequently confirmed by Grover et al. (2023) and UTMOST-NS (Dunn et al., 2023b). It was confidently detected by the online pipeline in the first post-glitch observation at UTMOST-NS, taken on January 1 2023 (MJD 59945), with $\ln K_1^* = 79$. It is similar to other glitches observed in this pulsar, with a moderate amplitude of $\Delta\nu_p/\nu = (328 \pm 1) \times 10^{-9}$ and a change in frequency derivative of $\Delta\dot{\nu}/\dot{\nu} = 1.2_{-0.6}^{+0.7} \times 10^{-3}$. No recovery is observed for this glitch. Zubieta et al. (2024) reported a recovering component with $\Delta\nu_d/\nu = 14.59(3) \times 10^{-9}$ and $\tau_d = 124(2)$ d. It is possible that we do not detect a recovering component is because we have only 179 d of post-glitch timing data

for this event, i.e. less than two e-foldings if the decay timescale reported by Zubieta et al. (2024) is accurate.

There are several previously reported glitches within our dataset for this pulsar. Two, reported as occurring at MJDs 57296.5 (Lower et al., 2020) and 57346 (Jankowski et al., 2015a; Lower et al., 2020; Basu et al., 2020), are small, with $\Delta\nu/\nu \lesssim 2 \times 10^{-9}$. We do not include these glitches in our timing model here. They are not detected by the HMM offline glitch search (note they lie below the 90% upper limit $\Delta\nu^{90\%}/\nu = 4.7 \times 10^{-9}$ for the relevant stretch of data; see Section 5.6.6) and visual inspection of the residuals does not reveal any obvious glitch signature. The other two glitches, at MJDs 57468 and 58241, are similar in size to the latest glitch and are clearly visible. We recover glitch parameters which are consistent with previous studies (Jankowski et al., 2016; Lower et al., 2020; Basu et al., 2020; Liu et al., 2024). We find moderate evidence for a recovery term in the glitch at MJD 58241 ($\ln \mathcal{B}_{\Delta\nu_d \neq 0}^{\Delta\nu_d = 0} = 5.9$), with $\Delta\nu_d/\nu = 5.3_{-5.3}^{+3.8}$ and $\Delta\tau_d = 56_{-56}^{+203}$. This recovery was not reported by other analyses (Lower et al., 2020; Basu et al., 2020; Liu et al., 2024).

PSR J1803–2137

This glitch is identified by eye during the initial timing process as occurring during the gap between the EW and NS datasets, and is a large glitch, with $\Delta\nu/\nu = 4592_{-93}^{+99} \times 10^{-9}$ and $\Delta\dot{\nu}/\dot{\nu} = 5.2_{-3.5}^{+3.0} \times 10^{-3}$. This glitch was previously reported by Keith et al. (2024) as occurring at MJD 58920 with $\Delta\nu_p/\nu = (4702 \pm 11) \times 10^{-9}$. This value is in significant tension with our measurement, which we ascribe to the glitch occurring in a 700-day observing gap. This glitch has also been recorded in the Jodrell Bank Observatory (JBO) glitch catalogue⁷ (Basu et al., 2022) as occurring at MJD 58958 ± 40 , with $\Delta\nu/\nu = (4960.6 \pm 1.3) \times 10^{-9}$ and $\Delta\dot{\nu}/\dot{\nu} = (8.50 \pm 0.3) \times 10^{-3}$ – different again from the values reported here and by Keith et al. (2024).

PSR J1836–1008

This glitch is identified by eye as occurring during the gap between the EW and NS datasets, and was subsequently confirmed in the offline glitch search with $\ln K_1^* = 2.1 \times 10^2$ with an estimated glitch epoch between MJD 59427 and MJD 59458. It is the second glitch reported in this pulsar. With an estimated $\Delta\nu/\nu = 36 \pm 3 \times 10^{-9}$ it is an order of magnitude larger than the previous known glitch, which had $\Delta\nu/\nu = 3.6(1) \times 10^{-9}$ (Basu et al., 2022). No significant change in $\Delta\dot{\nu}/\dot{\nu}$ is detected. It was also detected by Keith et al. (2024) as occurring at MJD 58950, with $\Delta\nu/\nu = (32.8 \pm 0.9) \times 10^{-9}$. Like previous authors, we find no evidence for a decaying component.

PSR J1902+0615

This glitch is identified by eye as occurring between MJD 58589 and MJD 58595, in April 2019. It was subsequently confirmed in the offline glitch search, with $\ln K_1^* = 7.2$. It has not been reported elsewhere. Unlike the other glitches detected in the combined dataset, this glitch did not occur in the observing gap – the UTMOST-EW data for this pulsar extends up to July 2020. However, it was not detected in previous searches (Lower et

⁷<https://www.jb.man.ac.uk/pulsar/glitches/gTable.html>

al., 2020; Dunn et al., 2022b) because the cut-off date for this pulsar in the data release analysed in those studies is November 2018. The glitch is small, with $\Delta\nu/\nu = (7.6 \pm 0.2) \times 10^{-10}$, but it is the largest of the seven known glitches in this pulsar. Other reported glitches have fractional sizes between 2×10^{-10} and 4×10^{-10} (Yuan et al., 2010; Espinoza et al., 2011; Basu et al., 2022).

5.6.6 | Glitch upper limits

The automated nature of the HMM glitch detector allows us to compute systematic upper limits on the size of glitches that could have been detected in a given analysis. This involves first creating fake datasets which mimic the real datasets, i.e. they have the same spin parameters, and white and red noise injected using LIBSTEMPO at the levels recovered by ENTERPRISE. Glitches of various sizes are injected into these fake datasets, and the detection probability as a function of glitch size, denoted $P_d(\Delta\nu)$, is estimated by running the HMM on the fake datasets and checking in each instance whether a glitch candidate is recovered. We take the benchmark detection probability to be 90%, and denote the glitch size at which 90% of injected glitches are successfully detected by $\Delta\nu^{90\%}$. Previous sensitivity estimates of the HMM have employed a binary search method to determine $\Delta\nu^{90\%}$ (Lower et al., 2021; Dunn et al., 2022b; Dunn et al., 2023a). In this work we adopt a technique borrowed from sensitivity estimation in the context of continuous gravitational wave searches (Banagiri et al., 2019; Abbott et al., 2021d; Abbott et al., 2022d) and assume the detection probability is a sigmoid in $\log_{10} \Delta\nu$ of the form

$$P_d(\Delta\nu) = [1 + \exp\{-r[\log_{10}(\Delta\nu/\text{Hz}) - c_0]\}]^{-1}. \quad (5.38)$$

We fit for the dimensionless parameters r and c_0 using the LMFIT package (Newville et al., 2014), and invert the best-fit sigmoid curve to find the value of $\Delta\nu^{90\%}$ satisfying $P_d(\Delta\nu^{90\%}) = 0.9$.

The above procedure applies to both the online and offline glitch searches. In the online glitch searches, we randomly choose 10 30-ToA windows within the total UTMOST-NS dataset and compute the value of $\Delta\nu^{90\%}$ for each. The injected glitch epochs are chosen uniformly within the window. We expect some variation in these values as the observing cadence and telescope sensitivity (and hence the ToA uncertainties) both vary by approximately a factor of two over the two years of operation (see Appendix 5.A1). For each pulsar we take the average over the set of chosen windows and quote a mean $\Delta\nu^{90\%}$ value. The distribution of $\Delta\nu^{90\%}/\nu$ is shown in the top panel Figure 5.8, with a mean $\Delta\nu^{90\%}/\nu$ of 4.7×10^{-8} . This is comparable to the sensitivity achieved in the offline search of the UTMOST-EW dataset conducted by Dunn et al. (2022b), where the mean $\Delta\nu^{90\%}/\nu$ value is 1.9×10^{-8} .

The upper limits in the offline search case are computed in the same way, except that the complete EW+NS datasets are used and the injected glitch epochs are chosen uniformly within the total timespan of each dataset. In cases where the EW+NS dataset for a pulsar is divided by the presence of large glitches, we derive separate upper limits for each section. The bottom panel of Figure 5.8 shows the distribution of $\Delta\nu^{90\%}/\nu$ for the offline searches. Compared to the equivalent distribution in Dunn et al. (2022b), as well as the upper limits for the online pipeline as shown in the top panel of the same figure, the distribution is shifted towards smaller values, with a mean of 5.3×10^{-9} compared

to 1.9×10^{-8} in Dunn et al. (2022b) and 4.7×10^{-8} for the online glitch searches. The minimum upper limit for the offline search is 2.4×10^{-12} , achieved for PSR J0437–4715, compared to the minimum 6.8×10^{-11} achieved for PSR J1939+2134 in the online case. Similarly the maximum upper limit in the offline search is 1.3×10^{-7} , achieved for PSR J1338–6204, compared to the maximum 5.0×10^{-7} achieved for PSR J1703–4851 in the online case.

In the online searches, the primary factors determining sensitivity are the typical ToA uncertainty and observing cadence (Melatos et al., 2020; Dunn et al., 2023a). As discussed in Section 5.6.3, for the online searches we adopt a fixed configuration of the HMM glitch detector for all pulsars, and do not incorporate information about the characteristics of timing noise in the pulsar. In contrast, in the offline searches, information about the timing noise is used (see Appendix 5.A3), so that the intrinsic properties of the pulsar also play a role in determining the sensitivity; see Section 6 of Melatos et al. (2020).

Small glitches have recently been reported in PSR J1048–5832 (Zubieta et al., 2023; Liu et al., 2023; Zubieta et al., 2024) and PSR J0835–4510 (Zubieta et al., 2024) which are within the timespan covered by the EW+NS dataset. The four reported glitches in PSR J1048–5832 have fractional amplitudes ranging between 0.91×10^{-9} and 9.9×10^{-9} , while the small glitch in PSR J0835–4510 reported by Zubieta et al. (2024) has a fractional amplitude of 2×10^{-10} . None of these glitches are detected by the offline HMM search. They are smaller than our 90% upper limits, which are 2.5×10^{-8} for PSR J1048–5832 and 1.4×10^{-8} for PSR J0835–4510.

5.7 | Conclusion

In total, the UTMOST-NS timing programme produced 269 pulsar-years of new timing baseline across 167 pulsars, and 3665 hours of integrated observing time. In this paper we describe the programme, and report some of the first scientific results drawing on the full sample of observed pulsars, with a focus on irregularities in the timing behaviour. We augment the new UTMOST-NS data with existing data from the earlier UTMOST programme using the east-west arm (Jankowski et al., 2019; Lower et al., 2020), and execute two Bayesian analysis pipelines: ENTERPRISE (Ellis et al., 2020) selects between models with different combinations of timing noise and non-zero second frequency derivative, and a HMM-based glitch detector (Melatos et al., 2020) searches for glitches and produces systematic upper limits on the size of undetected glitches for each object.

We investigate the timing noise properties across the population, and find results consistent with earlier studies (Shannon and Cordes, 2010; Parthasarathy et al., 2019; Lower et al., 2020) which suggest that the pulsar characteristic age $\tau_c \propto \nu \dot{\nu}^{-1}$ is correlated with the contribution of timing noise to the total timing residuals, σ_{RN} . Assuming $\sigma_{\text{RN}} \propto \nu^a \dot{\nu}^b$, we find $a = -0.92^{+0.31}_{-0.29}$ and $b = 0.80^{+0.12}_{-0.13}$. We also find that fixing the scaling of σ_{RN} with observing timespan according to the measured noise PSD spectral index does not improve our ability to estimate the contribution of timing noise to the timing residuals. The scatter in the relationship between σ_{RN} and its predicted value increases by approximately 10% when the scaling with observation timespan takes its theoretically expected form. Motivated by this result, we test the consistency of the estimated noise PSD parameters under variations in data set and methodology, and find that the esti-




Figure 5.8: Distribution of 90% confidence upper limits on fractional glitch size for the online (*top*) and offline (*bottom*) glitch searches described in Sections 5.6.3 and 5.6.4 (blue histograms), compared to the observed distribution of fractional glitch sizes recorded in the JBO glitch catalogue (orange histograms). All histograms are normalised so that they integrate to one. The online upper limits include one value per pulsar, 158 total. The offline upper limits include one value per time span analysed per pulsar. When a data set has been split into pre- and post-glitch sections or has been subdivided because of strong timing noise, there are multiple time spans per pulsar – in total there are 173 offline upper limits. The JBO values include all glitches with positive $\Delta\nu$, 676 in total. In the case of the online search, the mean upper limit across the sample is $\Delta\nu^{90\%}/\nu = 4.6 \times 10^{-8}$ with a maximum of 5.0×10^{-7} (PSR J1703–4851) and a minimum of 6.8×10^{-11} (PSR J1939+2134). In the case of the offline search, the mean upper limit is 5.3×10^{-9} , with a maximum of 1.3×10^{-7} (PSR J1709–4429) and a minimum of 2.4×10^{-12} (PSR J0437–4715).

mated parameters of the noise PSD sometimes change significantly as additional data are incorporated. There is a bias towards recovery of shallower spectral indices β when we analyse the combined EW+NS datasets compared to the EW datasets alone, using the same methodology (ENTERPRISE) in both cases. The mean shift in β is $\Delta\beta = -0.57 \pm 0.14$. No significant shift in the amplitude A_{red} is observed. Moreover, no significant shift in either A_{red} or β is observed when comparing the results of ENTERPRISE against TEMPONEST, using the same dataset.

We investigate the presence of secular $\dot{\nu}$ terms for the pulsars in our sample. Even when simultaneously modelling the timing noise and $\dot{\nu}$ term, measurements of $\dot{\nu} \neq 0$ are still contaminated by timing noise, as previously noted by other authors (Vargas and Melatos, 2023a; Keith and Nițu, 2023). In particular, the ensemble variation in the measured $\dot{\nu}$ can exceed the uncertainty reported by ENTERPRISE by a typical factor of approximately three when the spectral index of the timing noise PSD is approximately six. This effect is closely connected to the low-frequency cutoff in the timing noise model. Setting the cutoff at $1/T_{\text{span}}$, we recover nine non-zero $\dot{\nu}$ values, with the implied braking index ranging between $-(2.4 \pm 2.4) \times 10^4$ and $(5.0 \pm 0.7) \times 10^4$. If an additional timing noise model with the cutoff at $1/(2T_{\text{span}})$ is evaluated, the only pulsar with a significant measurement of $\dot{\nu}$ is the Crab pulsar, for which we measure $n = 2.54 \pm 0.03$. We expand on the work of Vargas and Melatos (2023a) to give explicit analytic conditions under which “anomalous” braking indices are expected (i.e. $|n| \gg 1$) in terms of the parameters estimated by ENTERPRISE. These conditions involve the amplitude and spectral index of the timing noise PSD, the spin frequency and its first derivative, and the observing timespan. We show that for all but two pulsars in our sample, variation in the measured n due to timing noise is expected theoretically to overwhelm the underlying secular braking index n_{pl} , under the assumption that $1 \lesssim n_{\text{pl}} \lesssim 7$. The median timing baseline required to overcome this effect is 280 yr for pulsars with spectral indices close to 4, and 3.9×10^7 yr for those with spectral indices close to 6.

Finally we report on the presence of glitches in our sample, using the HMM pipeline to perform both an “online” low-latency search for glitches, which was in place during active operation, and an “offline” search which was performed afterwards. We analyse 17 glitches in total, of which three were detected while the UTMOST-NS timing programme was active and one, in PSR J1902+0615, is previously unpublished. Glitch parameters are estimated using ENTERPRISE, and we check for the presence of an exponentially decaying term via Bayesian model selection. We find evidence for decaying components in seven glitches among four pulsars, with timescales ranging between 1.9 ± 0.7 d and $616_{-359.1}^{+359.3}$ d and recovery fractions on the order of 10^{-3} , within the typically observed ranges for large glitches (e.g. Yu et al. 2013; Lower et al. 2021; Liu et al. 2024). For all three large glitches in PSR J0835–4510 we measure a short decay timescale between 1 and 10 d, and a recovery fraction of approximately 6×10^{-3} , matching several previous studies (Flanagan, 1990; McCulloch et al., 1990; Dodson et al., 2002; Sarkissian et al., 2017; Zubieta et al., 2024). Measurement of these decay timescales, particularly those on the order of days, is enabled by the high cadence of the observing programme.

We also report estimates of the sensitivity of both the online and offline glitch detection pipelines, leveraging the fast and automated nature of the HMM pipeline to derive upper limits on the size of undetected glitches through the use of automated injection studies. We find that the online pipeline achieves an average sensitivity of

$\Delta\nu^{90\%}/\nu = 4.6 \times 10^{-8}$, while the offline pipeline achieves a significantly better average sensitivity of $\Delta\nu^{90\%}/\nu = 5.3 \times 10^{-9}$ due to the explicit incorporation of information about the timing noise.

This study of bright southern pulsars is the last major project to run at the Molonglo Telescope. It is being decommissioned in 2024 after almost 60 years of operation as a significant Australian radio astronomy facility.

Acknowledgements

Parts of this research are supported by the Australian Research Council (ARC) Centre of Excellence for Gravitational Wave Discovery (OzGrav) (project numbers CE170100004 and CE230100016). LD is supported by an Australian Government Research Training Program Scholarship. This work was performed in part on the OzSTAR national facility at Swinburne University of Technology. The OzSTAR program receives funding in part from the Astronomy National Collaborative Research Infrastructure Strategy allocation provided by the Australian Government. The Molonglo Observatory is owned and operated by the University of Sydney with support from the School of Physics and the University. The UTMOST project is also supported by the Swinburne University of Technology.

Data availability

The data products underlying this work will be made available through the Pulsar Data Portal (<https://pulsars.org.au>) and are also available upon reasonable request to the corresponding author.

5.A1 | Flux density calibration and telescope sensitivity for the survey period

We monitor the sensitivity of the telescope over the course of the timing programme using 12 bright “calibrator” pulsars which exhibit relatively steady emission properties, in particular a consistent flux density from observation to observation. We use these calibrator pulsars because the NS arm upgrade did not include e.g. noise diodes, which can be used to measure variations in the sensitivity of radio telescopes (e.g. Maron et al. 2013; Jankowski et al. 2018; Kumamoto et al. 2021). The modified radiometer equation (Keane, 2010) gives the proportionality between the flux density of the pulsar and the observed S/N in terms of a number of properties of the telescope and the pulsar:

$$S = \frac{T_{\text{sys}} + T_{\text{sky}}}{G \sqrt{n_p \Delta f t_{\text{int}}}} \sqrt{\frac{\delta}{1 - \delta}} (\text{S/N}), \quad (5.39)$$

where S is the flux density, G is the telescope gain in units of KJy^{-1} , T_{sys} is the system temperature in K, T_{sky} is the sky temperature in K, n_p is the number of polarisations, Δf




Figure 5.9: Sensitivity of the UTMOST-NS array as a function of time, normalised against the average sensitivity during the observations taken between MJD 59598 and MJD 59652, indicated by the dashed vertical lines.

is the observing bandwidth, t_{int} is the integration time, δ is the duty cycle of the pulsar⁸. Thus tracking S/N for these pulsars with approximately constant S gives a proxy for the variation in the SEFD, $G^{-1}T_{\text{sys}}$, over time, as all other parameters are either constant or well-known for each observation.

The sensitivity curve we produce is quoted in relative terms — we calculate a correction factor which adjusts the measured S/N values so that they are comparable to the S/N values obtained during a period of high sensitivity between MJDs 59598 and 59652. For each calibrator pulsar we find the mean S/N value within this calibration window, and then compute the ratio between this mean S/N and the S/N at each observation. This gives a point estimate of the telescope sensitivity, normalised against the window. All S/N values are scaled to a reference observing time of 300 s by multiplying by a factor of $\sqrt{300 \text{ s}/t_{\text{int}}}$. The final sensitivity curve is then calculated by taking the running mean over windows of 15 observations (encompassing all calibrator pulsars), and is shown in Figure 5.9.

We use the telescope sensitivity curve along with *a priori* knowledge of the flux densities of the sample of 12 pulsars to construct an approximate flux density scale for the pulsars in our sample. To do this we use estimates of the flux densities at 843 MHz (Jankowski et al. 2019; S. Johnston, priv. comm.) for the same twelve calibrator pulsars used to construct the sensitivity curve. The gain of the telescope is estimated to be 0.25 KJy^{-1} , and the system temperature approximately 170 K based on the S/N values observed from the calibrator pulsars, for an overall SEFD of 680 Jy (Mandlik et al., 2024). Given these values, along with (i) an estimate of the sky temperature for each

⁸As δ approaches 1 equation (5.39) becomes inaccurate, but the pulsars in the UTMOST-NS dataset have $0.005 < \delta < 0.27$. Note that highly scattered pulses may have $\delta > 1$ nominally yet are still detectable under some circumstances.

pulsar derived from the temperature maps of Haslam et al. (1982) scaled to our observing frequency and (ii) the telescope sensitivity curve, we estimate the flux density from UTMOST-NS data for every observation of each calibrator pulsar via equation (5.39), where S/N includes the correction factor derived from the telescope sensitivity curve. With the estimated system parameters quoted above, we find that the flux density estimates are in fair agreement with the Jankowski et al. (2019) values, which were calculated with reference to absolute flux density calibrated observations taken at the Parkes radio observatory (Jankowski et al., 2018). However, the UTMOST-NS flux densities tend to be somewhat lower than the Jankowski et al. (2019) values, and an overall correction factor of 1.07 is applied to the estimated SEFD to align the two sets of values, so that the mean ratio between the estimated flux densities from UTMOST-NS observations and the Jankowski et al. (2019) values is unity.

We thus obtain a flux density scale for all the pulsars in our sample, using the sensitivity curve of Figure 5.9 along with the telescope parameters and overall correction factor quoted in the previous paragraph. We do not use the derived flux densities further in this work. However they are available as one of the data products through the Pulsar Data Portal⁹ (Bailes et al., in prep.). We caution that these values are indicative only, and should be treated as accurate to within a factor of approximately two.

5.A2 | Influence of timing noise on the measured $\ddot{\nu}$ for $|\beta - 4| < 1$

This appendix presents the calculation of $\delta\ddot{\nu}$ when the spectral index of the timing noise PSD, β , satisfies $|\beta - 4| < 1$. The analysis is much the same as the $\beta \approx 6$ case discussed in Section 5.5.1, but we assume a Langevin equation of the form in equation (5.21) rather than (5.14) – i.e. there is a random walk in ν rather than $\dot{\nu}$. The strength of the white noise term $\xi(t)$ is parametrised by σ_ν , and the phase residual PSD is

$$P_\phi(f) = 2\sigma_\nu^2(2\pi f)^{-4}. \quad (5.40)$$

Again we must take into account the difference in units by multiplying the ENTERPRISE PSD by ν^2 , and account for the mismatch in spectral indices by constructing a $P^{\beta=4}(f)$ with $\beta = 4$ and A_{red} adjusted to match σ_{RN} values. The amplitude of this adjusted PSD is given by

$$A_{\text{red}}^{\beta=4} = \sqrt{\frac{3}{\beta - 1}} A_{\text{red}} \left(\frac{f_{\text{low}}}{f_{\text{yr}}} \right)^{-(\beta-4)/2}. \quad (5.41)$$

Then the value of σ_ν can again be estimated by equating $P_\phi(f)$ and $\nu^2 P^{\beta=4}(f)$:

$$\sigma_\nu = \sqrt{\frac{f_{\text{yr}} \nu^2 (2\pi)^4}{24\pi^2}} A_{\text{red}}^{\beta=4}. \quad (5.42)$$

The dispersion in ν over a timespan T is $\delta_\nu = \sigma_\nu T^{1/2}$, and the corresponding dispersion in $\ddot{\nu}$ is $\delta\ddot{\nu} = 2\delta_\nu/T^2$. This leads to

$$\delta\ddot{\nu} \sim 2\sqrt{\frac{f_{\text{yr}} \nu^2 (2\pi)^4}{24\pi^2}} A_{\text{red}}^{\beta=4} T^{-3/2}, \quad (5.43)$$

⁹<https://pulsars.org.au>

as in equation (5.22).

5.A3 | HMM setup for offline glitch searches

To completely specify the HMM used for a given pulsar in the offline analysis (Section 5.6.4), several parameters must be chosen, including the choice of timing noise model and the parameter controlling the strength of that timing noise, as well as the parameters controlling the DOI, which is the grid of $(\nu, \dot{\nu})$ values allowed in the HMM. Note that here $(\nu, \dot{\nu})$ specifies a deviation away from a fixed timing model (with parameters indicated by a subscript LS in the main text), and so for this application the DOI is always symmetric about zero in both ν and $\dot{\nu}$. For convenience we introduce the uncertainty in ν , denoted $\delta\nu$, and similarly for $\dot{\nu}$. These are the uncertainties due to the ToA measurement error – the uncertainty in ν across a ToA gap of length z_k with ToA measurement errors σ_k and σ_{k+1} (in units of cycles) is given by

$$\delta\nu_k = z_k^{-1}(\sigma_k^2 + \sigma_{k+1}^2)^{1/2}, \quad (5.44)$$

and similarly for $\dot{\nu}$ the uncertainty is given by

$$\delta\dot{\nu}_k = z_k^{-2}(\sigma_k^2 + \sigma_{k+1}^2)^{1/2}. \quad (5.45)$$

The characteristic uncertainties for a given pulsar, $\delta\nu$ and $\delta\dot{\nu}$ are taken to be the median values of $\delta\nu_k$ and $\delta\dot{\nu}_k$ over all ToA gaps. Finally, we denote that median ToA gap length for a given pulsar by \bar{z} .

As noted by Dunn et al. (2023a), the nature of the HMM means that a higher cadence of observations is not always advantageous if the aim is to detect the smallest glitches in the data. Smaller z reduces the dispersion in ν and $\dot{\nu}$ during each ToA gap, but it also increases the uncertainty in ν and $\dot{\nu}$ during the gap; see equations (5.44) and (5.45). It therefore makes sense in some cases to *remove* a subset of the data before analysis, in order to minimise

$$f(\bar{z}) = \sqrt{\delta\nu^2 + (\sigma\bar{z}^\alpha)^2}, \quad (5.46)$$

where $f(\bar{z})$ is a measure of the average spread in ν across ToA gaps in a given dataset, including the contribution from both the uncertainty due to ToA error bars ($\delta\nu$) and the timing noise model ($\sigma\bar{z}^\alpha$). The value of the exponent α depends on the timing noise model chosen; for RWF0 we have $\alpha = 1/2$ and for RWF1 we have $\alpha = 3/2$. For each dataset we find what \bar{z} minimises $f(\bar{z})$, and trim the dataset to remove gaps smaller than this value. We impose an upper bound on the adopted optimal \bar{z} of 30 d. When no timing noise is indicated by the ENTERPRISE analysis we always take the optimal \bar{z} to be 30 d.

When choosing the HMM parameters, there are three cases to cover: the case where a model with no timing noise is preferred, the case where a timing noise model with $2 < \beta < 5$ is preferred¹⁰, and the case where a timing noise model with $\beta > 5$ is preferred.

5.A3.1 | No timing noise

In the case where a model with no timing noise is preferred, we select the RWF1 timing noise model in the HMM, and set the timing noise strength parameter σ_j based on $\delta\nu$.

¹⁰The case $\beta < 2$ is excluded by our adopted priors; see Table 5.1.

Over a gap of length z , the dispersion in ν due to the random walk under RWF1 is $\sigma_{\dot{\nu}}z^{3/2}$, and so

$$\sigma_{\dot{\nu}} = \delta\nu\bar{z}^{-3/2}, \quad (5.47)$$

so that over the median ToA gap length \bar{z} the dispersion in ν is equal to the typical uncertainty in ν due to the ToA measurement uncertainty. The ν bounds on the DOI are taken to be $\nu_{\pm} = 10\delta\nu$, and the bin sizing to be $\epsilon_{\nu} = \delta\nu/2$. The $\dot{\nu}$ bounds are taken to be $\dot{\nu}_{\pm} = 2\nu_{\pm}/T_{\text{span}}$, and the bin spacing is $\epsilon_{\dot{\nu}} = 2\dot{\nu}_{\pm}/11$. The precise factors involved here are arbitrary but conservative – given that these pulsars do not exhibit detectable levels of timing noise, we can afford to sample ν and $\dot{\nu}$ finely, as we do not have to accommodate a wide range in either parameter.

5.A3.2 | Timing noise with $2 < \beta < 5$

In the case where a model including timing noise with $2 < \beta < 5$ is preferred, we select RWF0, which produces a PSD with $\beta = 4$. Equation (5.42) gives a recipe for setting σ_{ν} , the parameter which controls the strength of the random walk in the HMM, according to the value of A_{red} returned by ENTERPRISE. Under the RWF0 model, the dispersion in ν over a gap of length z is $\sigma_{\nu}z^{1/2}$. We take the DOI bounds on ν to be $\nu_{\pm} = \sigma_{\nu}T_{\text{span}}^{1/2}$. Similarly the DOI bounds on $\dot{\nu}$ are $\dot{\nu}_{\pm} = \sigma_{\nu}T_{\text{span}}^{-1/2}$. The bin spacing in ν is taken to be $\epsilon_{\nu} = \max(\sigma\bar{z}^{1/2}/2, \delta\nu/2)$. That is, the ν bins are spaced by either the typical dispersion in ν over a ToA gap due to the random walk, or by the typical uncertainty in ν due to ToA measurement error, whichever is larger. The bin spacing in $\dot{\nu}$ is taken to be $\epsilon_{\dot{\nu}} = \min(2\dot{\nu}_{\pm}/3, \delta\dot{\nu}/2)$. We aim to space the $\dot{\nu}$ bins according to the typical uncertainty in $\dot{\nu}$ due to ToA uncertainty, but in some cases where $\dot{\nu}_{\pm}$ is particularly small this leads to only a single $\dot{\nu}$ bin. We find that the HMM produces large numbers of clear false alarms in this case, so we enforce a minimum of at least three $\dot{\nu}$ bins.

5.A3.3 | Timing noise with $\beta > 5$

The final case is $\beta > 5$, and we select the RWF1 model for this case, which produces a PSD with $\beta = 6$. As with the RWF0 model discussed in Section 5.A3.2, we have a recipe [equation (5.19)] for setting the strength of the random walk, in this case denoted $\sigma_{\dot{\nu}}$, based on the value of A_{red} returned by ENTERPRISE. The DOI bounds on ν and $\dot{\nu}$ are set to $\nu_{\pm} = \sigma T_{\text{span}}^{3/2}$ and $\dot{\nu}_{\pm} = \sigma T_{\text{span}}^{1/2}$ respectively. The bin spacing in $\dot{\nu}$ is taken to be $\epsilon_{\dot{\nu}} = \max(\sigma\bar{z}^{1/2}/2, \delta\dot{\nu}/2)$. As for the the case of ν spacing in Section 5.A3.2, the $\dot{\nu}$ bins are spaced either according to the typical dispersion in $\dot{\nu}$ over a ToA gap or by the uncertainty in $\dot{\nu}$ due to ToA uncertainty, whichever is larger. In some cases the spacing in $\dot{\nu}$ and the initially chosen value of $\sigma_{\dot{\nu}}$ are such that transitions to neighbouring $\dot{\nu}$ bins are suppressed, and we find that this causes the HMM to produce large numbers of clear false alarms. A similar issue arises with $\dot{\nu}$ spacing in Section 5.A3.2. We include an adjustment to $\sigma_{\dot{\nu}}$ at this point, taking $\sigma_{\dot{\nu}} = \max(\sigma_{\dot{\nu}}, \epsilon_{\dot{\nu}}\bar{z}^{-1/2})$, so that $\sigma_{\dot{\nu}}$ is large enough to allow the hidden state to transition between neighbouring $\dot{\nu}$ bins. Finally we take the bin spacing in ν as $\epsilon_{\nu} = \max(\sigma\bar{z}^{3/2}, \delta\nu/2)$.

6

Search for continuous gravitational waves from neutron stars in five globular clusters with a phase-tracking hidden Markov model in the third LIGO observing run

This chapter has been submitted for publication to Physical Review D as L. Dunn et al., “Search for continuous gravitational waves from neutron stars in five globular clusters with a phase-tracking hidden Markov model in the third LIGO observing run”.

Abstract

A search is performed for continuous gravitational waves emitted by unknown neutron stars in five nearby globular clusters using data from the third Laser Interferometer Gravitational-Wave Observatory (LIGO) observing run, over the frequency range 100–800 Hz. The search uses a hidden Markov model to track both the frequency and phase of the continuous wave signal from one coherent segment to the next. It represents the first time that a phase-tracking hidden Markov model has been used in a LIGO search. After applying vetoes to reject candidates consistent with non-Gaussian artifacts, no significant candidates are detected. Estimates of the strain sensitivity at 95% confidence $h_{0,\text{eff}}^{95\%}$ and corresponding neutron star ellipticity $\epsilon^{95\%}$ are presented. The best strain sensitivity, $h_{0,\text{eff}}^{95\%} = 2.7 \times 10^{-26}$ at 211 Hz, is achieved for the cluster NGC6544.

6.1 | Introduction

Continuous gravitational waves (CWs) are long-lasting and nearly monochromatic signals which may be emitted by a variety of sources, but which have so far remained undetected (Riles, 2023; Wette, 2023). In the frequency band covered by existing terrestrial interferometers, one promising class of CW sources is rapidly rotating neutron stars located within the Milky Way. There are many possible emission mechanisms through which neutron stars can emit gravitational waves, involving both the crust and the interior. These include deformations created by thermal (Bildsten, 1998; Ushomirsky et al., 2000; Osborne and Jones, 2020; Hutchins and Jones, 2023), magnetic (Bonazzola and Gourgoulhon, 1996; Cutler, 2002; Payne and Melatos, 2004; Melatos and Payne, 2005; Payne and Melatos, 2006; Rossetto et al., 2023), and tectonic (Caplan and Horowitz, 2017; Gittins et al., 2021; Kerin and Melatos, 2022) effects, as well as fluid oscillations (particularly r -modes) (Andersson et al., 1999; Arras et al., 2003; Bondarescu et al., 2007). CW searches thus serve to probe the physics involved in these mechanisms.

Globular clusters are an interesting target for CW searches, as their dense cores provide a natural formation site for millisecond pulsars (MSPs), fast-spinning neutron stars which have been “recycled” by accretion from a binary companion (Katz, 1975; Clark, 1975; Alpar et al., 1982; Pooley et al., 2003). The gravitational wave strain scales as the square of the spin frequency for several important emission mechanisms (Riles, 2023). Once recycled, MSPs may also experience further accretion episodes later in their life, as the stellar density in the cluster core enhances the likelihood of stellar encounters which disrupt the orbits of debris disks (Wang et al., 2006; Abbott et al., 2017b; Jennings et al., 2020) or planets around the pulsar (Wolszczan and Frail, 1992; Bailes et al., 2011; Spiewak et al., 2018; Behrens et al., 2020; Nițu et al., 2022). Two CW searches targeting globular clusters specifically have been carried out in the past (Wette, 2023). Abbott et al. (2017b) carried out a search with 10 days of Initial Laser Interferometer Gravitational-Wave Observatory (LIGO) Science Run 6 data targeting unknown neutron stars in the globular cluster NGC 6544. Dergachev et al. (2019) carried out a search with Advanced LIGO Observing Run 1 data targeting a small sky region containing both the Galactic Center and the globular cluster Terzan 5. Searches directed at the Galactic Center are broadly similar to globular cluster searches, being wide parameter space searches targeting an interesting sky region, and a number of these searches have been performed previously (Aasi et al., 2013; Dergachev et al., 2019; Piccinni et al., 2020; Abbott et al., 2022f). Besides the two specific searches, globular cluster MSPs have been targeted in searches for known pulsars (Abbott et al., 2010; Aasi et al., 2014; Pitkin et al., 2015; Abbott et al., 2017e; Abbott et al., 2019d; Abbott et al., 2022f).

This paper presents a search for CW emission from unknown neutron stars in five globular clusters in the Milky Way. We allow for the possibility that the phase model of the signal does not follow a simple Taylor expansion but includes some additional stochasticity (called “spin wandering”). The search uses a hidden Markov model (HMM) to track both the frequency and phase of the wandering signal (Suvorova et al., 2016; Suvorova et al., 2018; Melatos et al., 2021). HMMs have been used in a number of CW searches in the past (Abbott et al., 2017c; Abbott et al., 2019c; Millhouse et al., 2020; Middleton et al., 2020; Sun et al., 2020a; Abbott et al., 2021d; Jones and Sun, 2021; Beniwal et al., 2021; Beniwal et al., 2022; Abbott et al., 2022b; Abbott et al., 2022g; Abbott et al.,

2022c; Vargas and Melatos, 2023b) but previously tracked only the frequency, typically using the \mathcal{F} -statistic for isolated targets (Jaranowski et al., 1998; Cutler and Schutz, 2005) or the \mathcal{J} -statistic for targets in binaries (Suvorova et al., 2017). In contrast, the HMM implementation used in this paper, developed by Melatos et al. (2021), employs a modified version of the \mathcal{B} -statistic (Prix and Krishnan, 2009; Melatos et al., 2021), which is a function of both frequency and phase. Phase tracking improves the sensitivity by down-weighting noise features, that are not approximately phase-coherent over the search.

The structure of the paper is as follows. In Section 6.2 we describe the HMM framework in general, its implementation in this specific search, and the modified \mathcal{B} -statistic which serves as the coherent detection statistic. In Section 6.3 we discuss the setup of the search, including target selection and the parameter domain. In Section 6.4 we discuss the candidates returned by the search, the veto procedures used to reject spurious candidates, and the results of applying the vetoes. In Section 6.5 we estimate the sensitivity of the search to CW emission from each cluster. In Section 6.6 we summarise the results and look to future work.

6.2 | Algorithms

In this section we briefly recap the HMM formalism and its algorithmic components. We define the signal model (Section 6.2.1), derive the form of the phase-sensitive \mathcal{B} -statistic which serves as the coherent detection statistic throughout this search (Section 6.2.2), and explain how the \mathcal{B} -statistic is combined with phase-frequency transition probabilities as part of a HMM (Section 6.2.3).

6.2.1 | Signal model

Following Jaranowski et al. (1998), the mass quadrupole gravitational wave signal from an isolated rotating neutron star emitting at twice the spin frequency can be written as a linear combination of four components

$$h(t) = A^\mu h_\mu(t), \quad (6.1)$$

with $\mu \in \{1, 2, 3, 4\}$, where we adopt the Einstein summation convention. The A^μ are the “amplitude parameters” given by Eqs. (32)–(35) in Jaranowski et al. (1998). They do not evolve and depend on the characteristic strain amplitude h_0 , the cosine of the inclination angle $\cos \iota$, the polarisation angle ψ and the initial signal phase Φ_0 .¹ The $h_\mu(t)$ are combinations of the antenna pattern functions [Eqs. (12) and (13) of Jaranowski et al. (1998)] and the phase of the signal $\Phi(t)$, whose evolution is expanded as a Taylor series, viz.

$$\Phi(t) = 2\pi \sum_{k=0}^s f^{(k)} \frac{t^{k+1}}{(k+1)!} + \Phi_w(t) + \Phi_{\text{SSB}}(t). \quad (6.2)$$

¹Strictly speaking the A^μ also depend on the angle between the total angular momentum of the star and the axis of symmetry, θ , as well as the angle between the interferometer arms, ζ (Jaranowski et al., 1998). The A^μ are multiplied by a factor of $\sin \zeta \sin^2 \theta$, but $\zeta = \pi/2$ for the LIGO detectors, and we assume $\theta = \pi/2$ also, so this factor is unity.

In Eq. (6.2), $f^{(k)}$ is the k -th time derivative of the frequency evaluated at $t = 0$, $\Phi_w(t)$ is an additional stochastic wandering term, the nature of which is discussed in detail in Section 6.2.3, and $\Phi_{\text{SSB}}(t)$ is the phase contribution from the transformation between the solar system barycentre (SSB) and the detector frame, cf. Eq. (14) of Jaranowski et al. (1998). For the form of $h_\mu(t)$ see Eq. (36) in Ref. (Jaranowski et al., 1998).

6.2.2 | Phase-dependent \mathcal{B} -statistic

The \mathcal{B} -statistic was introduced by Prix and Krishnan (2009) as a Bayesian alternative to the well-known \mathcal{F} -statistic (Jaranowski et al., 1998). The \mathcal{F} -statistic is obtained by analytically maximizing the likelihood function over A^μ and hence h_0 , $\cos \iota$, ψ and Φ_0 . The \mathcal{B} -statistic is obtained by instead *marginalising* over these parameters with a physically motivated prior which is uniform in h_0 , $\cos \iota$, ψ , and Φ_0 . Melatos et al. (2021) modified this approach to leave Φ_0 as a free parameter, so that it can be tracked by a HMM as discussed in Section 6.2.3. Here we briefly justify the form of the phase-dependent \mathcal{B} -statistic and discuss how to compute it in practice. The discussion draws on Refs. (Jaranowski et al., 1998; Prix and Krishnan, 2009; Melatos et al., 2021).

We begin with the log likelihood

$$\ln \Lambda = \langle x|h \rangle - \frac{1}{2} \langle h|h \rangle, \quad (6.3)$$

where $x(t)$ is the detector data and $h(t)$ is the signal as defined in Eqs. (6.1) and (6.2). The inner product is defined as

$$\langle x|y \rangle = 4\Re \int_0^\infty df \frac{x(f)y^*(f)}{S_h(f)} \quad (6.4)$$

$$\approx \frac{2}{S_h(f_0)} \int_0^T dt x(t)y(t), \quad (6.5)$$

where $S_h(f)$ is the one-sided power spectral density of the noise at a frequency f , T is the total length of the data, and the approximate equality holds for finite-length signals localized spectrally to a narrow band around f_0 in which $S_h(f) \approx S_h(f_0)$ can be taken as constant (which is a good approximation for the quasi-monochromatic signals of interest here (Jaranowski et al., 1998)). Defining $x_\mu = \langle x|h_\mu \rangle$ and $\langle h_\mu|h_\nu \rangle = \mathcal{M}_{\mu\nu}$ we write

$$\ln \Lambda = A^\mu x_\mu - \frac{1}{2} A^\mu \mathcal{M}_{\mu\nu} A^\nu. \quad (6.6)$$

The phase-dependent \mathcal{B} -statistic is obtained by marginalising over $\cos \iota$, ψ , and h_0 :

$$\mathcal{B} = \int_{-1}^1 d(\cos \iota) \int_{-\pi/4}^{\pi/4} d\psi \int_0^{h_0^{\max}} dh_0 p(\cos \iota, \psi, h_0) \Lambda. \quad (6.7)$$

In Eq. (6.7), $p(\cos \iota, \psi, h_0) = C/h_0^{\max}$ is the prior, where h_0^{\max} is a large but arbitrary maximum h_0 and C is an irrelevant normalising constant. Since $p(\cos \iota, \psi, h_0)$ is a constant and can be taken outside the integral it is dropped going forward.

We now substitute Eq. (6.6) into Eq. (6.7) and perform the triple integral. Writing $\mathcal{A}^\mu = h_0 \mathcal{A}'^\mu$, with \mathcal{A}'^μ independent of h_0 , and collapsing the triple integral $\int[\dots]$ for notational convenience, we obtain

$$\mathcal{B} = \int [\dots] \exp \left(h_0 \mathcal{A}'^\mu x_\mu - \frac{h_0^2}{2} \mathcal{A}'^\mu \mathcal{M}_{\mu\nu} \mathcal{A}'^\nu \right). \quad (6.8)$$

We use the following identity to perform the integral over h_0 :

$$\int_0^X dx \exp(Ax - Bx^2) = \left(\frac{\pi}{4B} \right)^{1/2} \exp \left(\frac{A^2}{4B} \right) \left[\operatorname{erf} \left(\frac{A}{2\sqrt{B}} \right) + \operatorname{erf} \left(\frac{2BX - A}{2\sqrt{B}} \right) \right]. \quad (6.9)$$

We also take h_0^{\max} , which is arbitrary, to be sufficiently large that the second $\operatorname{erf}(\dots)$ term approaches unity. The result is

$$\mathcal{B} = \int_{-1}^{-1} d(\cos \iota) \int_{-\pi/4}^{\pi/4} d\psi \left\{ \left(\frac{\pi}{2\mathcal{A}'^\mu \mathcal{M}_{\mu\nu} \mathcal{A}'^\nu} \right)^{1/2} \times \exp \left[\frac{(\mathcal{A}'^\mu x_\mu)^2}{2\mathcal{A}'^\mu \mathcal{M}_{\mu\nu} \mathcal{A}'^\nu} \right] \left[\operatorname{erf} \left(\frac{\mathcal{A}'^\mu x_\mu}{\sqrt{2\mathcal{A}'^\mu \mathcal{M}_{\mu\nu} \mathcal{A}'^\nu}} \right) + 1 \right] \right\}. \quad (6.10)$$

The ingredients of the integrand — \mathcal{A}'^μ , $\mathcal{M}_{\mu\nu}$ and x_μ — are all straightforward to compute using existing functionality in LALSUITE, the standard software suite for analysis of LIGO, Virgo, and KAGRA data (LIGO Scientific Collaboration et al., 2018). We use the existing graphics processing unit (GPU) framework (Dunn et al., 2022c) to compute the \mathcal{F} -statistic and hence x_μ , the most expensive step. We also integrate Eq. (6.10) using GPUs, running a simple Simpson integrator on a 11×11 grid of $\cos \iota$ and ψ values. We verify using synthetic data that the relative error in the value of the \mathcal{B} -statistic when computed using an 11×11 grid compared to a finer 501×501 grid is typically small, with a mean fractional error of 1% and a fractional error exceeding 10% occurring in 0.7% of the 6.4×10^5 \mathcal{B} -statistic values computed for the purposes of this test.

When computing \mathcal{B} we ignore $\Phi_w(t)$ in Eq. (6.2). By assumption this term is small compared to the other terms in Eq. (6.2) over each coherent segment during which \mathcal{B} is computed, amounting to $\lesssim 1$ rad. The stochastic wandering is tracked instead by the HMM, as discussed in Section 6.2.3.

6.2.3 | HMMs and the Viterbi algorithm

The number of signal templates needed to cover an astrophysically relevant parameter domain grows quickly for a phase-coherent search, as the observational timespan T_{obs} increases. Hence a number of “semicoherent” techniques for CW searches have been developed, which perform phase-coherent computations on subdivided segments of length T_{coh} and combine the segments without enforcing phase coherence between them. Most semicoherent methods completely discard the phase information between segments (see e.g. Refs. (Suvorova et al., 2016; Krishnan et al., 2004; Bayley et al., 2019; Astone et al., 2014)). Some, termed “loosely coherent”, retain some phase information but do not enforce perfect phase continuity (Dergachev, 2010; Dergachev, 2012; Dergachev

and Papa, 2023). While there is some loss of sensitivity inherent in this approach — the sensitivity of semicoherent searches scales as $(T_{\text{coh}}T_{\text{obs}})^{-1/4}$ (Krishnan et al., 2004; Wette, 2012), compared to $T_{\text{obs}}^{-1/2}$ for coherent searches — there is a large computational saving. Some semicoherent methods also offer flexibility in the face of the unexpected. For example, although fully coherent methods can accommodate deviations away from an assumed signal model through explicit searches over the parameters determining the phase evolution (Abbott et al., 2008b; Ashok et al., 2021; Abbott et al., 2022a), if there is a stochastic component to the phase evolution which operates on timescales shorter than the observing timespan then a semicoherent approach may be needed. A HMM solved by the Viterbi algorithm to find the most likely sequence of signal frequencies $f(t)$ is one such approach, which has been used in a number of CW searches in the past (Abbott et al., 2017c; Abbott et al., 2019c; Millhouse et al., 2020; Sun et al., 2020a; Middleton et al., 2020; Abbott et al., 2021d; Jones and Sun, 2021; Beniwal et al., 2021; Beniwal et al., 2022; Abbott et al., 2022b; Abbott et al., 2022g; Abbott et al., 2022c; Vargas and Melatos, 2023b). This approach was extended to include tracking of $\Phi(t)$ by Melatos et al. (2021), and is the subject of the present paper. In this section we review the basic framework of HMMs and the way it is put into practice for the present, loosely coherent search.

HMMs

HMMs model the evolution and observation of systems fulfilling two important criteria: a) the evolution of the internal state of the system is Markovian, and b) the internal state is hidden from the observer but may be probed probabilistically through measurement. This section briefly describes the formalism involved; for a more detailed review of HMMs see Rabiner (1989).

The hidden state of the system is assumed here to be discretised and bounded, so that each state is in the set $\mathcal{Q} = \{q_1, q_2, \dots, q_{N_Q}\}$, where N_Q is the total number of hidden states. Time is likewise discretised, with the HMM occupying the hidden state $q(t_i) \in \mathcal{Q}$ at each timestep $t_i \in \{t_0, t_1, t_2, \dots, t_{N_T}\}$, where t_0 is an “initial timestep” before any observations have been made. A sequence of possible hidden states occupied during a particular realisation of the HMM is denoted by $Q = [q(t_0), q(t_1), q(t_2), \dots, q(t_{N_T})]$. The dynamics of the system are captured by the transition matrix $A(q_j, q_i)$, whose entries are defined as

$$A(q_j, q_i) = \Pr[q(t_{n+1}) = q_j \mid q(t_n) = q_i]. \quad (6.11)$$

At each timestep we assume that the system is observed. For a given realisation we write the sequence of observations as $O = [o(t_1), o(t_2), \dots, o(t_{N_T})]$. The set of possible observations need not be finite, discretised, or bounded. All we require is that there exists an emission probability function $L(o, q_i)$ defined as

$$L(o, q_i) = \Pr[o(t_n) = o \mid q(t_n) = q_i], \quad (6.12)$$

i.e. the probability of observing o at the timestep t_n , if the hidden state q_i is occupied at t_n .

Finally we specify a belief about the state of the system at the beginning of the analysis, i.e.

$$\Pi(q_i) = \Pr[q(t_0) = q_i]. \quad (6.13)$$

In this paper we have no reason to impose anything other than a uniform prior, viz. $\Pi(q_i) = N_Q^{-1}$.

Given the above ingredients, it is straightforward to compute the probability of a hidden state sequence Q , given a sequence of observations O :

$$\Pr(Q | O) = \Pi[q(t_0)] \prod_{n=1}^{N_T} L[o(t_n), q(t_n)] A[q(t_n), q(t_{n-1})]. \quad (6.14)$$

The goal is to compute $\arg \max_Q \Pr(Q | O)$.

Phase tracking

In this paper, the full observation is divided into segments of length $T_{\text{coh}} = T_{\text{obs}}/N_T$, which are analysed coherently. Each coherent segment is a timestep in the HMM. The hidden states are pairs of frequencies and phases at the start of a segment, $q(t_n) = [f(t_n), \Phi(t_n)]$. The phase-dependent \mathcal{B} -statistic, defined in Section 6.2.2, plays the role of $\ln L(o, q_i)$. It then remains to specify the transition matrix. We follow Melatos et al. (2021) and adopt the following model for the evolution of the signal frequency f and phase Φ :

$$\frac{df}{dt} = \sigma \xi(t), \quad (6.15)$$

$$\frac{d\Phi}{dt} = f, \quad (6.16)$$

where $\xi(t)$ is a zero-mean white noise term, i.e.

$$\langle \xi(t) \rangle = 0 \quad (6.17)$$

$$\langle \xi(t) \xi(t') \rangle = \delta(t - t') \quad (6.18)$$

and σ parametrises the noise strength. The forward transition matrix $A[(f_j, \Phi_j), (f_i, \Phi_i)]$ is readily obtained by solving the Fokker-Planck equation corresponding to Eqs. (6.15)–(6.18). For practical reasons it is useful to also know the *backward* transition matrix

$$A^b(q_i, q_j) = \Pr[q(t_n) = q_i | q(t_{n+1}) = q_j], \quad (6.19)$$

which is similarly obtained by solving the backward Fokker-Planck equation for this system.

Details of the transition matrices are given in Appendix 6.A1, and an illustrative figure is shown in Fig. 6.1. The top panel of Fig. 6.1 shows a heatmap of $A(q_j, q_i)$ as a function of $\Delta\Phi = \Phi(t_{n+1}) - \Phi(t_n)$ and $\Delta f = f(t_{n+1}) - f(t_n)$. As shown in Melatos et al. (2021) and Appendix 6.A1, it is approximately a truncated Gaussian centred on $(\Delta\Phi, \Delta f) = (0, 0)$, with positive covariance between $\Delta\Phi$ and Δf . The bottom panel shows slices through the distribution at $\Delta f = -1, 0, +1$ bins. These slices are the distributions used to compute the transition probabilities in the HMM. The $\Delta f = \pm 1$ bin distributions are identical up to the location of the peak and each contain approximately 20% of the total probability, while the $\Delta f = 0$ bin contains the other 60%. The transitions probabilities are broad in phase. All three distributions have similar widths. From this figure we see

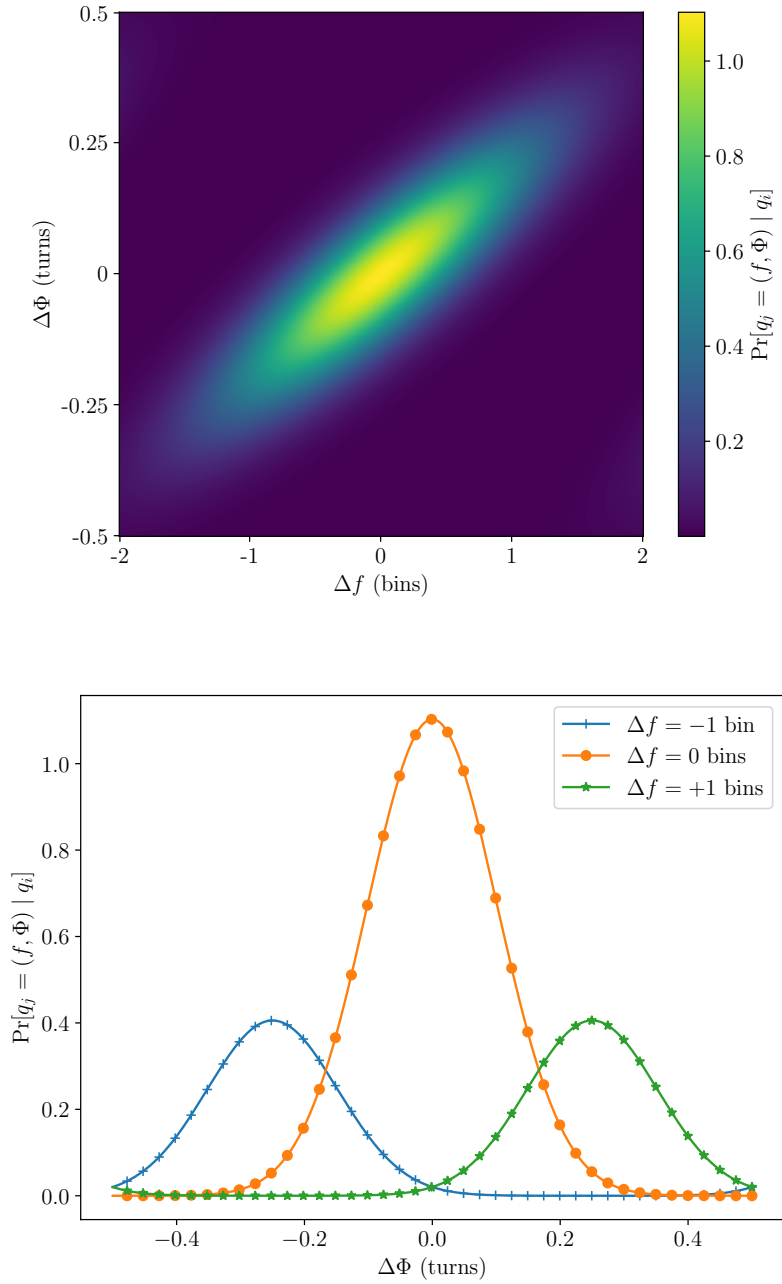


Figure 6.1: Illustrative plots of the transition matrix $A(q_j, q_i) = \Pr[q(t_{n+1}) = q_j \mid q(t_n) = q_i]$ according to the model described in Section 6.2.3. The top panel shows a heatmap of $A(q_j, q_i)$ where $q_i = [\Phi(t_n), f(t_n)]$ is an arbitrary initial state and $q_j = [\Phi(t_{n+1}), f(t_{n+1})]$ is the final state. The heatmap axes are $\Delta\Phi = \Phi(t_{n+1}) - \Phi(t_n)$ and $\Delta f = f(t_{n+1}) - f(t_n)$ in units of turns (cycles of 2π radians) and bins [of width $(2T_{\text{coh}})^{-1}$] respectively. The bottom panel shows slices of $A(q_j, q_i)$ as a function of $\Delta\Phi$ for $\Delta f = -1, 0, +1$ bins. Parameters: $T_{\text{coh}} = 5$ d, $\sigma = 1.76 \times 10^{-9} \text{ s}^{-3/2}$.

that the probability is concentrated in the same bin as the previous timestep, but some leakage into neighboring bins does occur.

Melatos et al. (2021) included a damping term $-\gamma f$ in Eq. (6.15). However, we are free to demand $\gamma < (2fT_{\text{coh}}^2)^{-1}$ and hence $|f(t_{n+1}) - f(t_n)| < (2T_{\text{coh}})^{-1}$, so that the damping does not cause $f(t)$ to move more than one bin per coherent segment. This renders the effect of $-\gamma f$ in Eq. (6.15) negligible, and in this work we do not include it. The approximation is justified in more detail in Appendix 6.A1.

Optimal Viterbi path

With the HMM fully specified, we turn to the question of how to use it in a CW search. To find signal candidates we use the Viterbi algorithm (Rabiner, 1989; Suvorova et al., 2016) to find the most likely sequence of hidden states,

$$Q^* = \arg \max_Q \Pr(Q | O). \quad (6.20)$$

More precisely, for each possible terminating frequency bin f_i we find the most likely sequence of hidden states which terminates in that bin,

$$Q^*(f_i) = \arg \max_Q \Pr[Q | O, f(t_{N_T}) = f_i]. \quad (6.21)$$

These paths are sorted according to their probability $\Pr[Q^*(f_i) | O]$. Paths exceeding a predetermined threshold are marked for follow-up analysis. The determination of the threshold as a function of the desired false alarm probability is discussed in Section 6.4.1. To maximize numerical accuracy, we work with log probabilities so that the products in Eq. (6.14) become sums. Hence we threshold candidates according to their log likelihood $\mathcal{L} = \ln \Pr[Q^*(f_i) | O]$.

To maximize computational performance the HMM tracking step is performed on GPUs, as with the computation of the \mathcal{B} -statistic described in Section 6.2.2.

6.3 | Setting up a globular cluster search

In this section we describe the astronomical targets and data inputs (Sections 6.3.1 and 6.3.2 respectively) for the globular cluster search performed in this paper, the choices of certain control parameters which dictate the behaviour of the HMM (Sections 6.3.3 and 6.3.4 respectively), and the astrophysical and computational motivations for the parameter domain and its associated grid spacing (Sections 6.3.5 and 6.3.6 respectively).

6.3.1 | Target selection

Globular clusters are dense stellar environments, with correspondingly high rates of dynamical interactions. They are prodigious manufacturers of low-mass X-ray binaries (LMXBs) (Katz, 1975; Clark, 1975; Pooley et al., 2003), which are the progenitors of MSPs, neutron stars which have been “recycled” to millisecond spin periods via accretion (Alpar et al., 1982). Although accretion from a binary companion offers several mechanisms to increase the mass quadrupole moment of a neutron star (Wagoner, 1984; Bildsten,

1998; Vigelius and Melatos, 2009), the search in this paper is not tuned to sources in binaries, as the \mathcal{B} -statistic in Section 6.2.2 does not include the effects of binary motion on the frequency modulation of the signal (doing so would require a blind search over the binary parameters, adding greatly to the cost of the search). Hence the ideal target in this paper is a neutron star that has spent some time in a binary system historically, so as to be recycled to spin frequencies in the LIGO band above 100 Hz (see Section 6.3.5), but was disrupted since by a secondary encounter and became an isolated source.

The catalog of Milky Way globular clusters compiled by Harris (1996) (2010 edition) contains parameters for 157 clusters. With limited resources, we are unable to search every known globular cluster for CW emission. Therefore we develop a simple ranking statistic to guide target selection. The rest of this section describes that statistic and the resulting target list.

Verbunt and Freire (2014) introduced a parameter² η which they refer to as the “single-binary encounter rate”. Roughly speaking it equals the reciprocal of the lifetime of a binary, before it is disrupted by a secondary encounter. It scales with globular cluster parameters as

$$\eta \propto \frac{\sqrt{\rho_c}}{r_c}, \quad (6.22)$$

where ρ_c is the core luminosity density and r_c is the core radius. Larger values of η are correlated with properties which are encouraging for the present search (Verbunt and Freire, 2014). First and foremost, isolated MSPs are more prevalent in high- η clusters, perhaps due to recycling followed by ejection from the binary by secondary encounters. Second, high- η clusters contain pulsars far from the core. A pulsar kicked out of the core by a dynamical interaction should sink back to the core due to dynamical friction, as long as it remains bound (Phinney, 1992). Hence detections of MSPs far from the core suggests that the secondary encounter which ejected the pulsar from the binary happened relatively recently so the overall rate of such encounters is indeed enhanced in high- η clusters. Finally, clusters with dense cores are observed to host pulsars with large period derivatives, which emit strongly in γ -rays, suggesting that these are energetic pulsars whose period derivative is intrinsic and not primarily due to acceleration in the cluster potential – examples include PSR J1823–3021A (Freire et al., 2011), PSR M28A (Vurgun et al., 2022), PSR J1835–3259B (Zhang et al., 2022), and possibly PSR J1740–5340B in the globular cluster NGC 6397, one of the clusters targeted in the present search (Zheng et al., 2024). At the time of writing, no pulsars of this kind are observed in globular clusters with low-density cores.

We follow Abbott et al. (2017b) and prioritize targets according to the cluster figure of merit $\eta^{1/2}D^{-1}$, where D is the distance to the cluster. The figure of merit follows from the age-based limit on strain amplitude previously used in searches for CWs from supernova remnants (Wette et al., 2008; Abbott et al., 2021d; Abbott et al., 2022d), viz.

$$h_0^{\max} = 1.26 \times 10^{-24} \left(\frac{3.3 \text{ kpc}}{D} \right) \sqrt{\frac{300 \text{ yr}}{a}}, \quad (6.23)$$

where a is the age of the object. Recalling that η is the reciprocal of the binary disruption timescale, and that shorter timescales are preferred for our application, we replace a

²Verbunt and Freire (2014) labelled this parameter γ , here we relabel it as η to avoid confusion with the γ introduced in Section 6.2.3.

by η^{-1} , following Ref. (Abbott et al., 2017b). This approach yields the figure of merit $\eta^{1/2}D^{-1} = \rho_c^{1/4}r_c^{-1/2}D^{-1}$ (up to a multiplicative constant which we ignore, as we are only interested in the ranking of the clusters).

Ranking the clusters listed in the Harris catalog (Harris, 1996) according to this figure of merit and taking the top five, we arrive at the list in Table 6.1. The cluster with the second highest figure of merit is NGC 6544, which was previously targeted in the 10-day coherent search using data from the sixth LIGO science run performed by Abbott et al. (2017b). We caution that the figure of merit is *not* directly related to an expected strain amplitude, despite its origin in the age-based limit for supernova remnants. It should instead be regarded as a convenient encapsulation of the two primary criteria for this search: isolated sources (higher η) are better, as are closer sources (lower D).

6.3.2 | Data

This paper analyzes data from the third observing run (O3) of the Advanced LIGO detectors (Abbott et al., 2023b), two dual-recycled Fabry-Pérot Michelson interferometers with 4 km arms located in Hanford, Washington (H1) and Livingston, Louisiana (L1) in the United States of America (Buikema et al., 2020). The data cover the time period 1st April 2019 15:00 UTC (GPS 1238166018) to 27th March 2020 17:00 UTC (GPS 1269363618), with a commissioning break between 1st October 2019 and 1st November 2019 (Goetz and Riles, 2023). We analyze a calibrated data stream (Sun et al., 2020b; Sun et al., 2021), and so are working with measurements of the dimensionless strain h .

The O3 LIGO data contain noise transients. A self-gating procedure mitigates their impact on continuous wave searches (Davis et al., 2021; Zweizig and Riles, 2021). The detector data are packaged into “Short Fourier Transforms” (SFTs) generated from 1800 s of data. Multiple SFTs are combined to construct the \mathcal{B} -statistic in each coherent segment of length T_{coh} (see Section 6.2.2). In the event that there are not enough data available to compute the \mathcal{B} -statistic in a given coherent segment, due to detector downtime, the value of $\ln L(o, q_i)$ in that segment is set to zero for all q_i . In this paper the only data gaps significant enough to prevent the computation of the \mathcal{B} -statistic for a coherent segment occur during the commissioning break mentioned above.

6.3.3 | Coherence time

The duration of each coherent step T_{coh} is an important parameter. It controls the sensitivity of the search, although the dependence is fairly weak, scaling roughly as $(T_{\text{obs}}T_{\text{coh}})^{-1/4}$ (Krishnan et al., 2004; Wette, 2012). It also plays a role in determining the HMM’s tolerance to deviations away from a deterministic signal model — as the difference between the true signal phase and the phase model approaches 1 cycle, decoherence occurs and significant signal-to-noise ratio is lost. This difference may arise either via stochastic spin wandering, viz. $\Phi_w(t)$ in Eq. (6.2), or via errors in the parameters which determine the first and third terms of Eq. (6.2). In this paper we cover a wide parameter space, so it makes sense to set T_{coh} relatively low to reduce the number of templates and hence the computational cost. In this search we adopt $T_{\text{coh}} = 5 \text{ d}$, c.f. $T_{\text{coh}} = 10 \text{ d}$ in some HMM searches published previously (Abbott et al., 2019c; Middleton et al., 2020; Abbott et al., 2022g; Abbott et al., 2022b).

| Name | Right ascension | Declination | r_c [pc] | $\log_{10}(\rho_c/L_\odot \text{pc}^{-3})$ | D [kpc] | $\eta^{1/2} D^{-1}$ |
|----------|--|--------------|----------------------|--|-----------|---------------------|
| NGC 6325 | 17 ^h 17 ^m 59.21 ^s | -23°45'57.6" | 6.8×10^{-2} | 5.52 | 7.8 | 11.8 |
| Terzan 6 | 17 ^h 50 ^m 46.38 ^s | -31°16'31.4" | 9.9×10^{-2} | 5.83 | 6.8 | 13.4 |
| NGC 6540 | 18 ^h 06 ^m 08.60 ^s | -27°45'55.0" | 4.6×10^{-2} | 5.85 | 5.3 | 25.4 |
| NGC 6544 | 18 ^h 07 ^m 20.58 ^s | -24°59'50.4" | 4.4×10^{-2} | 6.06 | 3.0 | 52.2 |
| NGC 6397 | 17 ^h 40 ^m 42.09 ^s | -53°40'27.6" | 3.3×10^{-2} | 5.76 | 2.3 | 65.5 |

Table 6.1: Cluster parameters for the five clusters targeted in this paper, from Harris (1996). The sky position (columns two and three) refers to the centroid of the light distribution. The core radius is denoted by r_c , the central luminosity density is denoted by ρ_c , and the distance is denoted by D . The final column is the figure of merit used to rank the clusters, as discussed in Section 6.3.1.

Tracking intrinsic spin wandering is a secondary priority in this paper. Electromagnetic measurements of timing noise in non-accreting pulsars, which are the objects primarily targeted in this search, suggest that $T_{\text{coh}} \sim 5$ d with the transition matrix described in Section 6.2.3 allows conservatively for more spin wandering than is observed, assuming that the spin wandering in the CW signal tracks the electromagnetic observations. Although there is significant variation among pulsars, the timescale of decoherence due to spin wandering is typically $\gtrsim 1$ yr (Jones, 2004). Hence any plausible choice of T_{coh} with regard to computational cost ($T_{\text{coh}} \lesssim 30$ d) allows for more spin wandering than what is expected astrophysically. As a precaution, we inject spin wandering above the astrophysically expected levels when assessing the performance of the HMM method in Sections 6.3.4 and 6.5. This allows us to verify that we remain sensitive to a wide range of signal models, including some that are unlikely *a priori*.

6.3.4 | HMM control parameters

Timing noise strength

The phase-tracking implementation of the HMM assumes that the evolution of the rotational phase is driven by a random walk in the frequency, with a strength parametrised by a control parameter σ [see Equations (6.15)–(6.18)]. Our approach here is to take T_{coh} as a free parameter which we adjust based on the particular details of the search, whether that involves prior belief about the nature of the spin wandering, or a desire to keep the number of sky position/spindown templates under control. With T_{coh} fixed, we then set σ to be consistent with the signal model implied by our choice of T_{coh} . Setting σ appropriately is a subtle task. If σ is too large, the transition probabilities become uniform in both frequency and phase, the dynamical constraints imposed by the noise model are lost, and the method reduces approximately to the phase-maximized \mathcal{F} -statistic (Jaranowski et al., 1998). If σ is too small then the HMM struggles to track stochastic spin wandering or other unexpected phase evolution. Melatos et al. (2021) recommended

$$\sigma = \sigma^* = (4T_{\text{coh}}^3)^{-1/2} \quad (6.24)$$

but did not validate Eq. (6.24) empirically.

There is no unique prescription for choosing σ . The uncertainty about the physical origin of the time-varying mass-quadrupole moment makes it difficult to argue in general that one choice of σ (reflecting the assumed nature of the underlying noise process) is optimal for a given search. However, σ endows the signal model with flexibility and allows the HMM to absorb small errors in parameters determining the secular spin evolution, e.g. sky position and spin-down rate. This in turn reduces the number of templates and hence the computational cost. As the HMM is restricted by construction to moving by at most one frequency bin per timestep, there is a limit to the parameter error which can be absorbed.³ The aim, then, is to check whether Eq. (6.24) strikes an acceptable balance, allowing for both secular and stochastic deviations away from a Taylor-expanded signal model, as reflected in our choice of T_{coh} , without discarding the advantage gained by phase tracking.

³For sufficiently large mismatches in sky position or spindown a significant fraction of the total power is also lost in each coherent step, whether or not the HMM is allowed to track the frequency evolution of the signal across coherent steps.

We test the suitability of Eq. (6.24) via synthetic data injections into Gaussian noise. For a range of h_0 values we compute the detection probability P_d for $\sigma = 0.2\sigma^*$, σ^* , and $5\sigma^*$ as well as for no phase tracking, at a fixed false alarm probability of 10^{-13} per frequency bin. We inject moderate frequency wandering via stepwise changes in \dot{f} ; in each of 52 7-day chunks of injected data \ddot{f} is randomly chosen from a uniform distribution spanning $-1.13 \times 10^{-18} \leq \ddot{f}/(1 \text{ Hz s}^{-2}) \leq 1.13 \times 10^{-18}$. The values of ϕ , f , and \dot{f} are updated self-consistently as the signal evolves. Note that the timescale of \dot{f} evolution is deliberately chosen not to coincide with $T_{\text{coh}} = 5 \text{ d}$ to challenge the algorithm. The sky position (right ascension α , declination δ) and initial values of f and \dot{f} are chosen at random. Signal recovery is performed at the injected sky position, with a 10^{-4} Hz frequency band centred on $f = f(0)$ and $\dot{f} = [f(T_{\text{obs}}) - f(0)]/T_{\text{obs}}$. The injection and recovery setup is specified in Table 6.2.

| Parameter | Value |
|--------------------------|--|
| α | $\mathcal{U}(0, 2\pi) \text{ rad}$ |
| $\cos \delta$ | $\mathcal{U}(-1, 1)$ |
| $f(0)$ | $\mathcal{U}(100, 700) \text{ Hz}$ |
| $\dot{f}(0)$ | $\mathcal{U}(-5 \times 10^{-10}, 0) \text{ Hz s}^{-1}$ |
| $\cos \iota$ | $\mathcal{U}(-1, 1)$ |
| ψ | 0.49 rad |
| ϕ_0 | 0 rad |
| Detectors | H1, L1 |
| $S_h(f)^{1/2}$ | $4 \times 10^{-24} \text{ Hz}^{-1/2}$ |
| Spin wandering timescale | 7 d |
| T_{coh} | 5 d |
| T_{obs} | 364 d |

Table 6.2: Injection and recovery parameters for the synthetic data tests used to motivate the choices of HMM control parameters as described in Section 6.3.4. $\mathcal{U}(a, b)$ denotes a uniform distribution between a and b .

In order to find the detection probability P_d at a given probability of false alarm P_{fa} we require an estimate of the noise-only distribution of the log likelihoods returned by the HMM. To obtain this we fit an exponential tail to the distribution of log likelihoods returned from a large number of pure Gaussian noise simulations, using the procedure outlined in Appendix 6.A2 and adopted in previous searches (Abbott et al., 2022g; Knee et al., 2024). The injected amplitude of the signal is characterised by $h_{0,\text{eff}}$, a combination of h_0 and $\cos \iota$ which determines approximately the signal-to-noise ratio for many CW search methods and is defined as (Messenger et al., 2015)

$$(h_{0,\text{eff}})^2 = h_0^2 \frac{[(1 + \cos^2 \iota)/2]^2 + (\cos \iota)^2}{2}. \quad (6.25)$$

For $\cos \iota = 1$ we have $h_{0,\text{eff}} = h_0$, for $\cos \iota = 0$ we have $h_{0,\text{eff}} = h_0/\sqrt{8}$, and averaging over a uniform distribution in $\cos \iota$ gives $h_{0,\text{eff}} = (2/5)^{1/2} h_0$. Figure 6.2 graphs P_d as a function of $h_{0,\text{eff}}$ for the choices of σ outlined above. The choice $\sigma = \sigma^*$ performs as well as or better than the alternatives as well as no phase tracking across the full $h_{0,\text{eff}}$ range. We adopt $\sigma = \sigma^*$ throughout the rest of this paper.




Figure 6.2: Detection probability P_d versus effective wave strain $h_{0,\text{eff}}$ for three choices of the HMM control parameter σ as well as a control experiment with zero phase tracking (see legend), as described in Section 6.3.4.

Number of phase bins

Another important tunable parameter in the HMM is the number of bins used to track the signal phase, N_ϕ . Melatos et al. (2021) fixed $N_\phi = 32$. It is worth checking how reducing N_ϕ affects performance. We seek to reduce N_ϕ without sacrificing sensitivity, because the computational cost is proportional to N_ϕ .

To this end we test $N_\phi = 4, 8,$ and 32 (c.f. Section 6.3.4). Sensitivity curves showing P_d as a function of $h_{0,\text{eff}}$ are displayed in Figure 6.3. The differences are relatively minor, so we proceed with $N_\phi = 8$ in what follows.

6.3.5 | Parameter ranges

Ranges must be set for four parameters: the frequency and frequency derivative of the CW signal, and the sky regions covered for each cluster. Here we discuss each in turn, and justify the choices made in the implementation of the search.

Frequency and frequency derivative

In principle the possible frequency of CW emission from neutron stars spans a wide range. Observed pulsar spin frequencies range from less than 1 Hz (Tan et al., 2018; Reig, 2011) up to 716 Hz (Hessels et al., 2006). Although the relationship between the rotation frequency f_{rot} and the CW signal frequency f depends on the precise nature of the emission mechanism, one normally assumes $f_{\text{rot}} \leq f \leq 2f_{\text{rot}}$ (Riles, 2023). We remind the reader that the notional target for this search is a recycled neutron star, whose rotation frequency is likely to fall between 0.1 kHz and 0.7 kHz, as for observed millisecond pulsars. Note that $\sim 80\%$ of pulsars above 0.1 kHz fall in the range $0.1 \leq f_{\text{rot}}/1 \text{ kHz} \leq 0.4$ (Manchester et al., 2005). There are also some practical considerations when it comes




Figure 6.3: Detection probability P_d versus effective wave strain $h_{0,\text{eff}}$ for three choices of the number of phase bins N_Φ as well as a control experiment with zero phase tracking, as described in Section 6.3.4. The three choices of N_Φ perform similarly.

to selecting the frequency range. The LIGO interferometers are polluted by Newtonian noise below ~ 20 Hz, and narrowband spectral features (instrumental lines) below ~ 0.1 kHz (Covas et al., 2018; Gravitational Wave Open Science Center, n.d.).

The spin-down rate of a CW signal is weakly constrained *a priori*. Under the optimistic assumption that CW emission is responsible for all of the spindown of the emitting body, the strain amplitude is equal to

$$h_0^{\text{sd}} = 2.6 \times 10^{-25} \left(\frac{1 \text{ kpc}}{D} \right) \left(\frac{I_{zz}}{10^{38} \text{ kg m}^2} \right)^{1/2} \times \left(\frac{100 \text{ Hz}}{f} \right)^{1/2} \left(\frac{|\dot{f}|}{10^{-11} \text{ Hz s}^{-1}} \right)^{1/2} \quad (6.26)$$

where I_{zz} is the component of the moment of inertia tensor about the rotation axis. The expected sensitivity of this search is roughly (Jaranowski et al., 1998; Wette et al., 2008; Sun et al., 2018)

$$h_0^{\text{sens}} \approx 35 S_h^{1/2} (T_{\text{obs}} T_{\text{coh}})^{-1/4}, \quad (6.27)$$

which implies $\dot{f} \sim -10^{-10} \text{ Hz s}^{-1}$ for $h_0^{\text{sens}} = h_0^{\text{sd}}$. To beat the spin-down limit we require $h_0^{\text{sens}} < h_0^{\text{sd}}$. We therefore wish to pick a bound on \dot{f} such that the cross-over point $h_0^{\text{sens}} = h_0^{\text{sd}}$ occurs at a high enough frequency that we beat the spin-down limit for a significant fraction of plausible signal frequencies. Figure 6.4 shows the h_0^{sens} curve with three h_0^{sd} curves overplotted with $-\dot{f} = 10^{-11} \text{ Hz s}^{-1}$, $-\dot{f} = 1 \times 10^{-10} \text{ Hz s}^{-1}$, and $-\dot{f} = 5 \times 10^{-10} \text{ Hz s}^{-1}$. We set $D = 5 \text{ kpc}$ in all cases. The $\dot{f} = -5 \times 10^{-10} \text{ Hz s}^{-1}$ curve crosses over h_0^{sens} at approximately $f = 0.8 \text{ kHz}$. Hence for $0.1 \leq f/1 \text{ kHz} \leq 0.8$ and $-5 \leq \dot{f}/10^{-10} \text{ Hz s}^{-1} \leq 0$ we expect to approach or beat the spin-down limit across the full frequency range for each cluster. The range includes the bulk of the astrophysical




Figure 6.4: Comparison of h_0^{sens} [Eq. (6.27)] with h_0^{sd} [Eq. (6.26)], with h_0^{sd} computed for three values of \dot{f} ($D = 5$ kpc in all cases). For the $\dot{f} = -5 \times 10^{-10} \text{ Hz s}^{-1}$ case, corresponding to the bound adopted in this paper (see Section 6.3.5), we have $h_0^{\text{sens}} < h_0^{\text{sd}}$ for $f \lesssim 0.8$ kHz.

prior on the signal frequency, recalling that the majority of MSPs have $f_{\text{rot}} < 0.4$ kHz and thus $f < 0.8$ kHz. We therefore adopt $0.1 \leq f/1 \text{ kHz} \leq 0.8$ and $-5 \times 10^{-10} \leq \dot{f}/1 \text{ Hz s}^{-1} \leq 0$ in the present paper. To facilitate data management, the total frequency range is divided into 0.5 Hz sub-bands which are searched independently.

Sky position

Four of the clusters targeted in this search (NGC 6325, NGC 6397, NGC 6544, and Terzan 6) are core-collapsed; their density profiles peak sharply at the core. At the same time, our cluster selection criterion favors clusters which produce disrupted binaries where the neutron star may be ejected from the core. Indeed clusters similar to those selected are observed to contain millisecond pulsars many core radii away from the centre of the cluster (Verbunt and Freire, 2014). Therefore, as a precaution, we extend the search to cover sky positions outside the core. Specifically, we extend the search out to the tidal radius as recorded in Harris (1996), ensuring that every point within this radius is covered by at least one sky position template.

6.3.6 | Parameter spacing

In this paper we search a fairly wide parameter domain (precisely how wide is discussed in Section 6.3.5). In general, multiple templates are needed to cover the domain. In this section we discuss how closely spaced the templates must be. Sophisticated work has been done on this problem in the context of CW searches (e.g. (Prix, 2007b; Prix, 2007a; Pletsch, 2010; Wette, 2015)), but we proceed empirically here, because the Viterbi tracking step makes it difficult to apply previous results.




Figure 6.5: Example mismatch profile $m(\dot{f})$ versus \dot{f} , as defined by Equation (6.28). The profile is generated from a single random realisation of both noise and signal, with all parameters other than \dot{f} fixed to their injected values. The injected \dot{f} value is denoted by \dot{f}_{inj} .

Frequency and frequency derivative

We follow previous HMM-based searches and take the frequency resolution to be $(2T_{\text{coh}})^{-1}$. Other choices are possible, but this choice has been successfully used many times previously (see e.g. Refs. (Suvorova et al., 2016; Abbott et al., 2017c; Abbott et al., 2021d; Abbott et al., 2022g; Abbott et al., 2022b)) and is assumed in Eq. (6.24).

The HMM accommodates some displacement between the \dot{f} value used to compute the \mathcal{B} -statistic and the true \dot{f} by letting the signal wander into adjacent frequency bins at the boundaries between coherent segments. To determine the \dot{f} resolution we empirically estimate the mismatch, i.e. the fractional loss of signal power which is incurred as the value of \dot{f} used in the search is displaced away from the true value, \dot{f}_{inj} , with

$$m(\dot{f}) = \frac{\mathcal{L}'(\dot{f}) - \mathcal{L}'(\dot{f}_{\text{inj}})}{\mathcal{L}'(\dot{f}_{\text{inj}})}, \quad (6.28)$$

where we define $\mathcal{L}' = \mathcal{L} - \mathcal{L}_{\text{noise}}$ and $\mathcal{L}_{\text{noise}}$ is the mean likelihood returned in noise over a 0.5 Hz sub-band. Note that $\mathcal{L}'(\dot{f})$ is the maximum \mathcal{L} value over a 0.5 Hz sub-band centred on the injected frequency, and so $m(\dot{f})$ does not approach 1 but plateaus at approximately 0.8, due to the maximisation over noise-only frequency paths (whereas $\mathcal{L}_{\text{noise}}$ is the mean noise-only \mathcal{L} value in a 0.5 Hz sub-band, not the maximum). We compute $m(\dot{f})$ by injecting moderately loud signals into Gaussian noise, with $h_0 = 6 \times 10^{-26}$ and $S_h^{1/2} = 4 \times 10^{-24} \text{ Hz}^{-1/2}$. An example of $m(\dot{f})$ is displayed in Fig. 6.5, generated from a single realisation of noise and signal. All mismatch curves are smoothed using a Savitzky-Golay filter. The shape of $m(\dot{f})$ is consistent across noise realisations. Hence we take the \dot{f} resolution to be the median difference between the two \dot{f} values satisfying

$m(\dot{f}) = 0.2$, drawn from 100 noise realisations. This gives an \dot{f} resolution of $\delta\dot{f} = 3.7 \times 10^{-12} \text{ Hz s}^{-1}$. We verify that the resolution is independent of the injected f and \dot{f} values and adopt it throughout the search.

Sky location

A mismatch in sky location induces a small apparent variation in the signal frequency due to a residual Doppler shift. Although the locations of globular clusters are well-known, they are extended objects on the sky. In this search we take a conservative approach and search out to the tidal radius of each cluster, as discussed in Section 6.3.5. In general a single sky pointing does not suffice to cover the full sky area of each cluster. As with a mismatch in \dot{f} , we expect the transition matrix in Equation (6.15) to track some of the residual Doppler variation. Nevertheless it is important to check how much variation can be accommodated before a significant amount of signal power is lost.

We specify the sky location by right ascension (α) and declination (δ), referenced to the J2000 epoch. We empirically estimate the mismatch profiles $m(\alpha)$ and $m(\delta)$ and thence derive resolutions in α and δ . As with \dot{f} , we smooth the mismatch profiles using a Savitzky-Golay filter. The treatment is somewhat more involved than in Section 6.3.6, as $m(\alpha)$ and $m(\delta)$ depend on the injected values α_{inj} and δ_{inj} as well as f , i.e. their shape is not a universal function of sky position offset, because the Doppler modulation varies across the sky. In particular, sky positions at declinations near the ecliptic poles experience relatively little Doppler modulation, and hence proportionally smaller residual Doppler modulations (Abbott et al., 2017b). In this paper, we target five relatively small regions of the sky so we do not produce a full skymap of $m(\alpha, \delta)$. Instead we measure the α and δ resolutions independently for each cluster. The clusters themselves are fairly small (diameter $\lesssim 10$ arcmin), so we approximate the α and δ resolutions to be uniform within each cluster.

Each cluster is covered by ellipses arranged in a hexagonal grid, with their semimajor and semiminor axes given by the resolutions in δ and α corresponding to $m(\delta) \leq 0.2$ and $m(\alpha) \leq 0.2$. The centre of one ellipse coincides with the cluster core, and the whole cluster out to the tidal radius is covered by at least one ellipse. The centres of the ellipses are taken as the sky position templates in the search. For simplicity we do not consider degeneracy between sky position and \dot{f} , which may reduce modestly the number of templates, but adds to the complexity of implementing the search.

Examples of $m(\alpha)$ and $m(\delta)$ are shown in the top two panels of Fig. 6.6, calculated for Terzan 6 at a frequency of 600 Hz. They are qualitatively similar, with well-defined troughs, but the resolutions, again defined as the distances between the two abscissae where $m = 0.2$, differ significantly – the resolution inferred from the $m(\alpha)$ profile is 1.9 arcmin, compared to 15 arcmin for $m(\delta)$. The bottom panel of 6.6 shows the measured resolution in α and δ as a function of signal frequency for a series of injections at the sky location of Terzan 6. Clearly there is a strong frequency dependence, and it would be wasteful to simply apply the finest resolution across a wide band. Consequently, we produce curves like those in the bottom panel of Fig. 6.6 for each target. We interpolate by fitting a power-law function $Af^k + B$ for A , k , and B using the LMFIT package (Newville et al., 2016) to find the resolution in each 0.5 Hz sub-band.




Figure 6.6: Example sky resolution computations. The left and middle panels show example mismatch profiles $m(\alpha)$ and $m(\delta)$ respectively, calculated for the sky location of Terzan 6 at a frequency of 600 Hz for a single random realisation of noise and signal. The resolutions are 1.9 arcmin and 15 arcmin — note the difference in scale along the horizontal axes. The right panel shows the resolution in right ascension (blue curve) and declination (orange curve) as a function of frequency for the sky location of Terzan 6. Each point is derived by averaging the mismatch measured from 100 random realisations, and the curves are interpolated from the injection results by fitting a power law function, as described in the text.

6.4 | Candidates

In this section we present the results of the search. We set the log likelihood threshold in terms of the desired false alarm probability P_{fa} in Section 6.4.1. We study the candidates produced by a first-pass analysis in Section 6.4.2 and apply vetoes to reject those not of astrophysical origin. Finally we study the veto survivors in more depth in Section 6.4.3.

6.4.1 | Thresholds

We set a threshold log likelihood \mathcal{L}_{th} above which candidates are kept for further analysis. The threshold is determined by the desired P_{fa} , and is computed according to the procedure outlined in Appendix 6.A2. In short, we obtain a sample of \mathcal{L} values calculated in the absence of a signal, and fit an exponential distribution to the tail of this sample. The exponential distribution is used to calculate the value of \mathcal{L}_{th} that gives the desired P_{fa} . Empirical noise-only distributions of \mathcal{L} are obtained from 10^4 random sky pointings in each of three 0.5 Hz bands beginning at 263.5 Hz, 322 Hz, and 703 Hz, chosen arbitrarily but checked to ensure that the detector amplitude spectral densities (ASDs) show no obvious spectral artifacts. We find that \mathcal{L}_{th} is similar across the three bands (within ± 1 of the mean value per cluster), so we take the mean value over the three bands to be \mathcal{L}_{th} as a good approximation. We do not obtain separate noise-only distributions for each cluster.

The choice of P_{fa} is subjective. Historically different searches have used a range of values (Abbott et al., 2017b; Abbott et al., 2019d; Piccinni et al., 2020; Abbott et al., 2021d; Abbott et al., 2022g; Abbott et al., 2022b). In this paper, we take a relatively aggressive approach (i.e. P_{fa} is relatively high), setting \mathcal{L}_{th} on a per-cluster basis such that across all templates in that cluster, we expect one false alarm on average, and so expect the full search to produce five false alarms. This total number of expected false alarms is broadly comparable to other published HMM-based searches, which have adopted P_{fa} values giving false alarm rates ranging from 0.15 per search (Abbott et al., 2021d) to 18 per search (Abbott et al., 2022g). The search for CWs from NGC 6544 carried out by Abbott et al. (2017b) adopted a false alarm rate of 184 per search, while the search targeting Terzan 5 and the galactic center carried out by Dergachev et al. (2019) did not set a threshold explicitly in terms of a false alarm rate.

The total number of sky templates and hence \mathcal{L}_{th} differs between clusters. The per-cluster sky template counts and values of \mathcal{L}_{th} corresponding to a false alarm rate of one per cluster are quoted in Table 6.3.

6.4.2 | First pass candidates and vetoes

The number of above-threshold candidates in each cluster is listed in the pre-veto column of Table 6.4. The total across all clusters is 6.4×10^7 , with the candidates distributed approximately evenly among the five clusters. As discussed in Section 6.4.2, these pre-veto numbers are lower bounds, as most candidates are from heavily disturbed sub-bands and are not saved. The candidate counts in the subsequent columns are complete.

While 6.4×10^7 is far greater than the five expected based on the arguments in Section 6.4.1, we emphasise that those arguments are based on the assumption that the noise is

| Cluster | Sky templates | \mathcal{L}_{th} |
|----------|---------------|---------------------------|
| NGC 6544 | 3022 | 278.9 |
| NGC 6325 | 8416 | 283.9 |
| NGC 6540 | 8957 | 284.2 |
| Terzan 6 | 28462 | 289.8 |
| NGC 6397 | 43909 | 291.9 |

Table 6.3: Number of sky templates and \mathcal{L}_{th} values for each target cluster as discussed in Section 6.4.1, chosen to give a false alarm rate of one per cluster on average. Targets are ordered by the number of sky templates.

approximately Gaussian. All but a handful at most of the candidates in the pre-veto column of Table 6.4 are due to non-Gaussian features in the data. In this section and the next we apply a sequence of veto procedures designed to reject candidates that are consistent with non-Gaussian noise.

Disturbed sub-band veto

Some of the 0.5 Hz sub-bands are so disturbed by non-Gaussian artifacts that they are unusable in the absence of noise cancellation (Tiwari et al., 2015; Davis et al., 2019; Driggers et al., 2019; Kimpson et al., 2024). In view of this, we impose a blanket veto on sub-bands which produce output files larger than 100 MB (corresponding to more than 1.67×10^5 candidates per file). Between 3.9% and 5.4% of the full search band is removed thus, with the smallest number of removed sub-bands in NGC 6544 (55 out of 1400), and the highest in NGC 6397 (75 out of 1400). These percentages are roughly in line with the all-sky O3a search reported by Abbott et al. (2021a), for which 13% of the observing band was removed due to an unmanageable number of candidates (noting that the band in that search extended to 2000 Hz). After this veto, approximately 4.9×10^5 candidates remain.

Known-lines veto

The known-lines veto is simple and efficient, rejecting many candidates with relatively little computational cost. The LIGO Scientific Collaboration maintains a list of known instrumental lines (Goetz et al., 2021; Gravitational Wave Open Science Center, n.d.). When assessing coincidence we must take into account that the line frequencies and bandwidths in Ref. (Goetz et al., 2021) are quoted in the detector frame, while the HMM frequency path of a candidate is recorded in the Solar System barycentre frame. The resulting Doppler shift amounts to approximately $\pm 10^{-4} f$ (Jones et al., 2022). For each candidate we check whether any portion of the HMM frequency path overlaps with a known line in either detector, after correcting for the Doppler shift. In the event of any overlap the candidate is vetoed. After this veto, 18252 candidates remain.

6.4.3 | Second pass survivors and vetoes

The 18252 survivors from the first pass vetoes are still daunting to follow up manually. We apply three more vetoes to further exclude candidates that are consistent with non-

Gaussian noise artifacts.

Cross-cluster veto

The cross-cluster veto excludes candidates which have frequencies overlapping (up to Doppler modulation) with candidates in another cluster. Candidates which appear at the same frequency in multiple clusters are consistent with non-Gaussian noise artifacts. The targeted false alarm rate of one per cluster makes it unlikely that, out of a pair of candidates at the same frequency but in different clusters, one is due to an astrophysical signal while the other is due to a Gaussian false alarm. Thus we veto any candidate for which there is a candidate in another cluster with respective frequencies f_1 and f_2 satisfying $|f_1 - f_2| < 10^{-4} f_1$. 17977 candidates are vetoed this way. We find that all of the vetoed candidates have at least 4 cross-cluster counterparts, i.e. distinct combinations of α , δ , and \dot{f} with at least one frequency path with $\mathcal{L} > \mathcal{L}_{\text{th}}$ overlapping with the vetoed candidate – multiple cross-cluster counterparts may be associated with the same cluster. After this veto, 275 candidates remain.

Single-interferometer veto

If a candidate is due to a non-Gaussian feature which is present in the data from only one interferometer, we expect the significance of the candidate to increase when data from only that interferometer is used. We thus veto candidates satisfying $\mathcal{L}_X > \mathcal{L}_U$, where \mathcal{L}_X is the log likelihood of the candidate template using only data from detector X and \mathcal{L}_U is the joint-detector log likelihood. Before carrying out this veto we group candidates that have frequencies within 10^{-2} Hz of each other, in order to reduce the computation required and enable manual verification of the results of the veto procedure. There are 10 candidate groups in total, appearing in two clusters, NGC 6397 and NGC 6544. To compute \mathcal{L}_X over the relevant frequency band we use the loudest sky position and \dot{f} template in each candidate group, and \mathcal{L}_U is taken to be the loudest log likelihood value in the candidate group. Plots showing \mathcal{L}_X in comparison to \mathcal{L}_U for each candidate are collected in Appendix 6.A3. In all but one case, one \mathcal{L}_X value peaks above \mathcal{L}_U , indicating that the candidate group is likely to be due to a disturbance in the corresponding interferometer and is not astrophysical. These candidate groups are vetoed.

After the single-interferometer veto, two candidates remain.

Unknown lines veto

The line list used for the known-lines veto (Goetz et al., 2021) includes only lines for which there is good evidence of a non-astrophysical origin. Other line-like features are visible by eye in plots of the detector ASDs but have not yet been confidently associated with a terrestrial source.

As a further check on the final remaining candidate group, we manually inspect the ASD in each detector to check whether the candidate group overlaps with a loud non-Gaussian feature in the data, taking into account the Doppler modulation. Candidates which overlap with a clear feature in an ASD are vetoed, on the grounds that we expect any astrophysical signal to be at most marginally visible in the detector ASDs at distances $\gtrsim 2$ kpc (Jaume et al., 2024). A plot showing the detector ASDs and corresponding \mathcal{L}




Figure 6.7: Individual detector ASDs (top panel) and joint-detector log likelihoods (bottom panel) for the candidate group which survived the single-interferometer veto. The vertical dotted lines in the top panel indicate the extent of the Doppler modulation, and the horizontal dashed line in the bottom panel indicates \mathcal{L}_{th} . The cluster name and approximate ending frequency of the loudest candidate within the group are shown at the top of the figure. There is a clear disturbance in the H1 data within the Doppler modulation window, and no corresponding feature in the L1 data. This candidate group is vetoed.

values for the loudest candidate in the remaining candidate group, in the cluster NGC 6397 at $f \approx 652.756$ Hz, is shown in Figure 6.7. There is a clear, broad peak in the H1 ASD within the Doppler modulation range of the candidate, with no corresponding feature visible in the L1 ASD. We therefore veto this candidate group.

After the unknown lines veto, no candidates remain.

For completeness, in Appendix 6.A4 we present plots analogous to Fig. 6.7 for all ten candidate groups which survive the cross-cluster veto.

6.4.4 | Summary of vetoes

After applying the veto procedures described in Sections 6.4.2 and 6.4.3, we are left with no surviving candidates. A summary of the results of the veto procedures is given in Table 6.4. The disturbed sub-band and known-lines vetos reject most of the candidates but leave 18252, still a significant number. The cross-cluster veto further rejects all but 275 candidates, which are grouped into ten narrow frequency bands. We find that for nine of the ten groups, one of the single-detector log likelihoods peaks at a higher value than the joint-detector peak, indicating that these candidates are due to a non-Gaussian feature which is peculiar to one detector. Hence these groups are vetoed. The remaining group is in the vicinity of a broad, loud disturbance in the H1 detector and is therefore also vetoed.

| Cluster | Pre-veto | DS | KL | CC | 1IFO | UL |
|----------|-------------------|--------|-------|-----|------|----|
| NGC 6544 | 9.3×10^6 | 162509 | 2918 | 86 | 0 | 0 |
| NGC 6325 | 1.1×10^7 | 68338 | 3354 | 0 | 0 | 0 |
| NGC 6540 | 1.1×10^7 | 68656 | 235 | 0 | 0 | 0 |
| Terzan 6 | 9.9×10^6 | 145782 | 9952 | 0 | 0 | 0 |
| NGC 6397 | 1.2×10^7 | 43566 | 1793 | 189 | 2 | 0 |
| Total | 6.4×10^7 | 488851 | 18252 | 275 | 2 | 0 |

Table 6.4: Results of applying the veto procedures described in Sections 6.4.2 and 6.4.3. Each column gives the number of candidates surviving after the corresponding veto. The abbreviated column headings are: disturbed sub-band (DS), known-lines (KL), cross-cluster (CC), single-interferometer (1IFO) and unknown lines (UL). Note that the number of pre-veto candidates is a lower bound, as the candidate list from any sub-band which produces more than 1.67×10^5 candidates is not saved and is instead marked as vetoed by the disturbed sub-band veto (see Section 6.4.2).

6.5 | Sensitivity

In this section we estimate the sensitivity of the search to CW emission from each targeted globular cluster. Specifically, for each cluster we estimate $h_{0,\text{eff}}^{95\%}$, the effective strain amplitude [defined in Eq. (6.25)] at which we would expect to detect a signal 95% of the time. We draw the distinction between the search sensitivity, defined as above, and an upper limit. An upper limit is a limit on the loudest astrophysical signal actually present in the data, and is necessarily estimated with reference to the results of the search (i.e. estimation of the upper limit in a given sub-band depends on the loudest candidate in that sub-band). By contrast the search sensitivity is not a function of the results of the search, only the search configuration. We prefer to estimate the search sensitivity here as it is straightforward to interpret and requires significantly less computation. It can be estimated in a small number of sub-bands and interpolated across the full search band, as described below.

In order to estimate the search sensitivity we inject synthetic signals into the real O3 data. We inject at sky positions slightly offset in right ascension from the true cluster position, so that the sky templates do not overlap with the search, so as to avoid the possibility of contamination by an astrophysical signal from the cluster. In a given 0.5 Hz sub-band we inject signals at ten evenly spaced $h_{0,\text{eff}}$ levels between 1.7×10^{-26} and 3.5×10^{-26} and determine the detection rate as a function of $h_{0,\text{eff}}$. We then fit these rates as a function of $h_{0,\text{eff}}$ to a sigmoid curve, following Abbott et al. (2021d). We use the best-fit parameters to determine the $h_{0,\text{eff}}$ value yielding $P_d = 0.95$, in that sub-band.

The off-target injections and recoveries are treated much the same as in the actual search. The key difference is that the centre of the injection-containing mock cluster is offset from the true cluster centre by five tidal radii in right ascension. The placement of α , δ , and \dot{f} templates mimics that of the real search, but to save computational resources we only search \dot{f} templates within $\pm 3\delta\dot{f}$ of the injected \dot{f} . We also search a reduced frequency range $\pm 5 \times 10^{-4}$ Hz around the injection frequency. The value of \mathcal{L}_{th} is the

| Cluster | $\overline{\mathcal{D}}^{95\%}$ |
|----------|---------------------------------|
| Terzan 6 | 116.3 |
| NGC 6397 | 117.0 |
| NGC 6325 | 119.3 |
| NGC 6540 | 120.2 |
| NGC 6544 | 125.9 |

Table 6.5: Estimated sensitivity depths $\overline{\mathcal{D}}^{95\%}$ for the five targeted clusters. The value of $\overline{\mathcal{D}}^{95\%}$ is calculated by taking the mean of the $\mathcal{D}^{95\%}$ values [see Eq. (6.29)] computed in six sub-bands.

same as in the real search⁴, and we claim a successful recovery if the maximum recovered value of \mathcal{L} exceeds \mathcal{L}_{th} regardless of which template returns the maximum. As we set \mathcal{L}_{th} on a per-cluster basis (see Section 6.4.1), we expect different sensitivity curves for each cluster, with clusters requiring fewer sky position templates (and hence smaller \mathcal{L}_{th} values) reaching deeper sensitivities.

We determine the value of $h_{0,\text{eff}}^{95\%}$ directly following the above recipe in a small number of 0.5 Hz bands (beginning at 168.2 Hz, 242.6 Hz, 424.3 Hz, 562.8 Hz, 622.2 Hz, and 702.0 Hz) and then interpolate to cover the full observing band. The value of $h_{0,\text{eff}}^{95\%}$ is a strong function of frequency, but to a good approximation it is directly proportional to the detector-averaged ASD $S_h(f)^{1/2}$. Therefore in order to interpolate the estimated $h_{0,\text{eff}}^{95\%}$ values in the six sub-bands to a $h_{0,\text{eff}}^{95\%}$ curve which covers the entire observing band we estimate the sensitivity depth $\mathcal{D}^{95\%}$, which is the constant of proportionality between $h_{0,\text{eff}}^{95\%}$ and $S_h(f)^{1/2}$ and is defined as

$$\mathcal{D}^{95\%} = \frac{S_h(f)^{1/2}}{h_{0,\text{eff}}^{95\%}}. \quad (6.29)$$

We compute $S_h(f)$ using `LALPULSAR_COMPUTEPSD`, averaging over all data used in the search from both detectors. There is some scatter in the computed $\mathcal{D}^{95\%}$ values in the six 0.5Hz sub-bands, with $0.96 < \mathcal{D}^{95\%} / \overline{\mathcal{D}}^{95\%} < 1.04$ for each cluster, where $\overline{\mathcal{D}}^{95\%}$ is the mean value of $\mathcal{D}^{95\%}$ across the six sub-bands. There is no clear trend in $\mathcal{D}^{95\%}$ as a function of frequency. As $\mathcal{D}^{95\%}$ is defined to be the constant of proportionality between $S_h(f)^{1/2}$ and $h_{0,\text{eff}}^{95\%}$, to obtain a curve of $h_{0,\text{eff}}^{95\%}$ for each cluster over the full band we divide $S_h(f)^{1/2}$ by the $\overline{\mathcal{D}}^{95\%}$ value for that cluster.

The resulting $h_{0,\text{eff}}^{95\%}$ sensitivity curves are shown in the top panel of Figure 6.8. The $\mathcal{D}^{95\%}$ values used to produce the $h_{0,\text{eff}}^{95\%}$ curves are listed in Table 6.5. The differences in sensitivity between the five clusters are due to the varying \mathcal{L}_{th} values for each cluster (see Table 6.3), which follow from the number of sky position templates (see Sections 6.3.5 and 6.3.6). The variation between clusters is fairly small: the least sensitive cluster (Terzan 6) is approximately 10% worse than the most sensitive (NGC 6544), achieving

⁴In the searches around these mock injections the number of templates is smaller than the number used in real searches, implying a smaller \mathcal{L}_{th} for the targeted false alarm rate. Nonetheless the aim is to ascertain the sensitivity of the real searches, and so we use the \mathcal{L}_{th} values quoted in Table 6.3.




Figure 6.8: Sensitivity estimates for the five targeted clusters. *Top panel*: characteristic wave strain at 95% detection probability, $h_{0,\text{eff}}^{95\%}$, as a function of signal frequency, f (in Hz), interpolated from the average estimated sensitivity depth defined in Eq. (6.29) over six 0.5 Hz bands. The curves for NGC 6397 and Terzan 6 are too close to be distinguished, as are the curves for NGC 6540 and NGC 6325. *Bottom panel*: minimum ellipticity $\epsilon^{95\%}$ as a function of f for the values of D listed in Table 6.1 and the fiducial value of I_{zz} in Eq. (6.30). The clusters are color-coded according to the legend.

$h_{0,\text{eff}}^{95\%} = 2.9 \times 10^{-26}$ and 2.7×10^{-26} respectively at the most sensitive frequency of $f = 211$ Hz.

Out of astrophysical interest, we also convert $h_{0,\text{eff}}^{95\%}$ in Figure 6.8 into estimates of the minimum ellipticity of a source which would have been detected by this search assuming emission at twice the rotational frequency, according to the relation (Beniwal et al., 2022)

$$\epsilon^{95\%} = 9.5 \times 10^{-6} \left(\frac{h_{0,\text{eff}}^{95\%}}{10^{-25}} \right) \left(\frac{I_{zz}}{10^{38} \text{ kg m}^2} \right)^{-1} \times \left(\frac{f}{100 \text{ Hz}} \right)^{-2} \left(\frac{D}{1 \text{ kpc}} \right). \quad (6.30)$$

Curves of $\epsilon^{95\%}$ versus signal frequency for the five clusters are shown in the bottom panel of Figure 6.8. The smallest minimum ellipticity estimates are achieved at the top of the observing band, 800 Hz, due to the f^{-2} factor in Eq. (6.30). The factor of D introduces larger variation between the clusters compared to the variation in $h_0^{95\%}$ values, and modifies the relative ordering – the smallest minimum ellipticity estimate is 1.6×10^{-7} , for the nearest cluster NGC 6397 ($D = 2.3$ kpc), while the largest is 5.2×10^{-7} for the farthest cluster NGC 6325 ($D = 7.8$ kpc).

6.6 | Conclusion

We carry out a search for continuous gravitational radiation from unknown neutron stars in five globular clusters. This is the first time that all but one of the chosen clusters have been targeted. For NGC 6544 which was previously searched using LIGO S6 data, the sensitivity is improved approximately one order of magnitude (Abbott et al., 2017b) (although note that the signal models and parameter domain differ, so a direct, like-for-like comparison is impossible). This is also the first search to use a phase-tracking HMM in conjunction with a phase-dependent version of the \mathcal{B} -statistic (Prix and Krishnan, 2009; Melatos et al., 2021); previous HMM searches tracked frequency only, not phase.

We find no significant candidates. We estimate the 95% effective strain amplitude sensitivity for each cluster, along with a corresponding estimate of the 95% ellipticity sensitivity. The minimum per-cluster values of $h_{0,\text{eff}}^{95\%}$ lie in the range $2.5 \leq h_{0,\text{eff}}^{95\%}/10^{-26} \leq 3.0$, achieved at approximately 211 Hz, while the minimum values of $\epsilon^{95\%}$ are achieved at the top end of the search band (800 Hz), and lie in the range $1 \leq \epsilon^{95\%}/10^{-7} \leq 6$. The lowest $h_{0,\text{eff}}^{95\%}$ values are obtained for the cluster NGC 6544, while the lowest $\epsilon^{95\%}$ values are obtained for NGC 6397.

We surpass the spin-down limit for $|\dot{f}| \gtrsim 10^{-11} \text{ Hz s}^{-1}$, with variation between clusters due primarily to the difference in distance. MSPs known from electromagnetic measurements have spin-down rates which are orders of magnitude below this limit (Manchester et al., 2005). Their spin down is typically attributed to magnetic dipole braking, with the recycling process suppressing the magnetic dipole moment (Konar and Bhattacharya, 1997; Melatos and Phinney, 2001; Melatos and Payne, 2005; Priymak et al., 2011; Mukherjee, 2017). The objects to which this search is sensitive are therefore quite unlike the known MSP population, but are instead likely to be “gravitars” – neutron stars whose spin-down behaviour is dominated by gravitational rather than electromagnetic

torques (Palomba, 2005). The inferred limits on the ellipticity of the star probe up to an order of magnitude below theoretical upper limits, viz. $\epsilon \lesssim 10^{-6}$ (Johnson-McDaniel and Owen, 2013; Morales and Horowitz, 2022).

Acknowledgments

Parts of this research are supported by the Australian Research Council (ARC) Centre of Excellence for Gravitational Wave Discovery (OzGrav) (project numbers CE170100004 and CE230100016) and ARC Discovery Project DP170103625. This material is based upon work supported by NSF's LIGO Laboratory which is a major facility fully funded by the National Science Foundation. LD is supported by an Australian Government Research Training Program Scholarship. LS is supported by the ARC Discovery Early Career Researcher Award, Project Number DE240100206. This work was performed in part on the OzSTAR national facility at Swinburne University of Technology. The OzSTAR program receives funding in part from the Astronomy National Collaborative Research Infrastructure Strategy allocation provided by the Australian Government.

6.A1 | Transition probabilities

The transition probabilities used in this work are essentially the same as those derived by Melatos et al. (2021). They are derived from the system of equations (6.15) and (6.16) except that we discard the damping term $-\gamma f$ included in Eq. (6.15) by Melatos et al. (2021). In this appendix we reproduce the form for $A^b(q_i, q_j) = \Pr[q(t_n) = q_i | q(t_{n+1}) = q_j] = p^B(t_n, q_i)$ which appears in Melatos et al. (2021), and subsequently justify the choice to drop the $-\gamma f$ term. We then write down the transition probabilities in the $\gamma \rightarrow 0$ limit, which we adopt in the search.

With $\Delta q = (\Delta f, \Delta \Phi)$ specifying the difference between the final and initial states, the distribution $p^B(t_n, \Delta q)$ is a wrapped Gaussian (Suvorova et al., 2018; Melatos et al., 2021),

$$p^B(t_n, \Delta q) = (2\pi)^{-1} (\det \Sigma)^{-1/2} \times \sum_{m=-\infty}^{\infty} \exp[-(\Delta q - Q_m) \Sigma^{-1} (\Delta q - Q_m)^T], \quad (6.31)$$

where we have

$$(Q_m)_1 = 0 \quad (6.32)$$

$$(Q_m)_2 = -m, \quad (6.33)$$

$$\Sigma_{11} = \frac{\sigma^2}{2\gamma} [1 - \exp(-2\gamma\tau)], \quad (6.34)$$

$$\Sigma_{12} = \Sigma_{21} = \frac{\sigma^2}{2\gamma^2} [1 - \exp(-\gamma\tau)]^2, \quad (6.35)$$

$$\Sigma_{22} = \frac{\sigma^2}{2\gamma^3} \left\{ 1 + 2\gamma\tau - [2 - \exp(-\gamma\tau)]^2 \right\}, \quad (6.36)$$

with $\tau = t_{n+1} - t_n = T_{\text{coh}}$. The condition $\gamma < (2fT_{\text{coh}}^2)^{-1}$ in Ref. (Melatos et al., 2021) ensures that the argument of the exponent, $\gamma\tau \sim (2fT_{\text{coh}})^{-1}$, is $O(10^{-6})$. Taylor expanding the exponentials in Eqs. (6.34)–(6.36) about $\gamma\tau = 0$ allows us to write these expressions as a prefactor multiplied by a power series in $\gamma\tau$ with a constant term that is $O(1)$, e.g.

$$\Sigma_{12} = \frac{\sigma^2\tau^2}{2} (1 - \gamma\tau + \dots). \quad (6.37)$$

Given $\gamma\tau \sim 10^{-6}$ we conclude by inspecting (6.37) and its counterparts that the effect of the damping term on the moments of the transition probability distribution is negligible. We therefore take the $\gamma \rightarrow 0$ limit and write

$$\Sigma_{11} = \sigma^2\tau, \quad (6.38)$$

$$\Sigma_{12} = \Sigma_{21} = \frac{\sigma^2\tau^2}{2}, \quad (6.39)$$

$$\Sigma_{22} = \frac{\sigma^2\tau^3}{2}. \quad (6.40)$$

The transition probabilities are obtained by substituting (6.38)–(6.40) into (6.31).

6.A2 | Detection thresholds

In order to reduce the number of candidates from a search to a manageable level, we set thresholds on the log likelihood of the paths returned from the search and do not retain any paths which fall below the threshold for followup analysis. The thresholds are set to target a desired false alarm probability P_{fa} , i.e. one in every P_{fa}^{-1} paths has $\mathcal{L} > \mathcal{L}_{\text{th}}$ in sub-bands where the data are not significantly polluted by non-Gaussian features.

Since in general we want P_{fa} to be small, setting \mathcal{L}_{th} requires knowledge of the behaviour of the distribution of path log likelihoods in the tail. In this work we assume that the tail of the \mathcal{L} distribution is exponential, as confirmed empirically in previous work (Abbott et al., 2022g; Knee et al., 2024). We use off-target searches in the real data to infer the parameters of the exponential distribution. We follow the method described in Appendix A 1 of Abbott et al. (2022g), which we briefly recapitulate here.

We assume that the likelihood distribution $p(\mathcal{L})$ has the form

$$p(\mathcal{L}) = A\lambda \exp[-\lambda(\mathcal{L} - \mathcal{L}_{\text{tail}})] \quad \text{for } \mathcal{L} > \mathcal{L}_{\text{tail}}, \quad (6.41)$$

where A and λ are constants and $\mathcal{L}_{\text{tail}}$ is an empirically determined cut-off above which the distribution is approximately exponential. We have a sample of N paths, of which N_{tail} have $\mathcal{L} > \mathcal{L}_{\text{tail}}$. The sample is generated by searching random sky positions and \dot{f} values in 0.5 Hz sub-bands that appear free from disturbance based on the detector noise ASDs. Then A is simply N_{tail}/N , and the maximum-likelihood estimate of λ is

$$\hat{\lambda} = \frac{N_{\text{tail}}}{\sum_{i=1}^{N_{\text{tail}}} \mathcal{L}_i - \mathcal{L}_{\text{tail}}}. \quad (6.42)$$

For a desired P_{fa} we then estimate \mathcal{L}_{th} to be

$$\mathcal{L}_{\text{th}}(P_{\text{fa}}) = -\frac{1}{\hat{\lambda}} \ln\left(\frac{NP_{\text{fa}}}{N_{\text{tail}}}\right) + \mathcal{L}_{\text{tail}}. \quad (6.43)$$

6.A3 | Single-interferometer veto

This appendix collects the plots of the single-detector log likelihood values \mathcal{L}_X in the vicinity of the ten candidate groups which pass through the single-interferometer veto described in Section 6.4.3. The plots are displayed in Figure 6.9. The log likelihoods are computed using the template of the loudest candidate in each group. The frequency of each candidate is indicated by a vertical dashed line, and the joint-detector log likelihood \mathcal{L}_J is indicated by a horizontal dashed line. A candidate group is vetoed if $\mathcal{L}_X > \mathcal{L}_J$ for either detector. Nine of the ten candidate groups are vetoed this way.

6.A4 | Unknown lines

This appendix collects plots of the single-detector ASDs in the vicinity of the ten candidate groups which survive the cross-cluster veto (Section 6.4.3). The plots are displayed in Figure 6.11. Although all but one of these candidate groups are vetoed by the single-interferometer veto (Section 6.4.3), we present the single-detector ASDs for completeness. Seven of the ten candidate groups show a disturbance in one ASD within the Doppler modulation window of the candidate group.

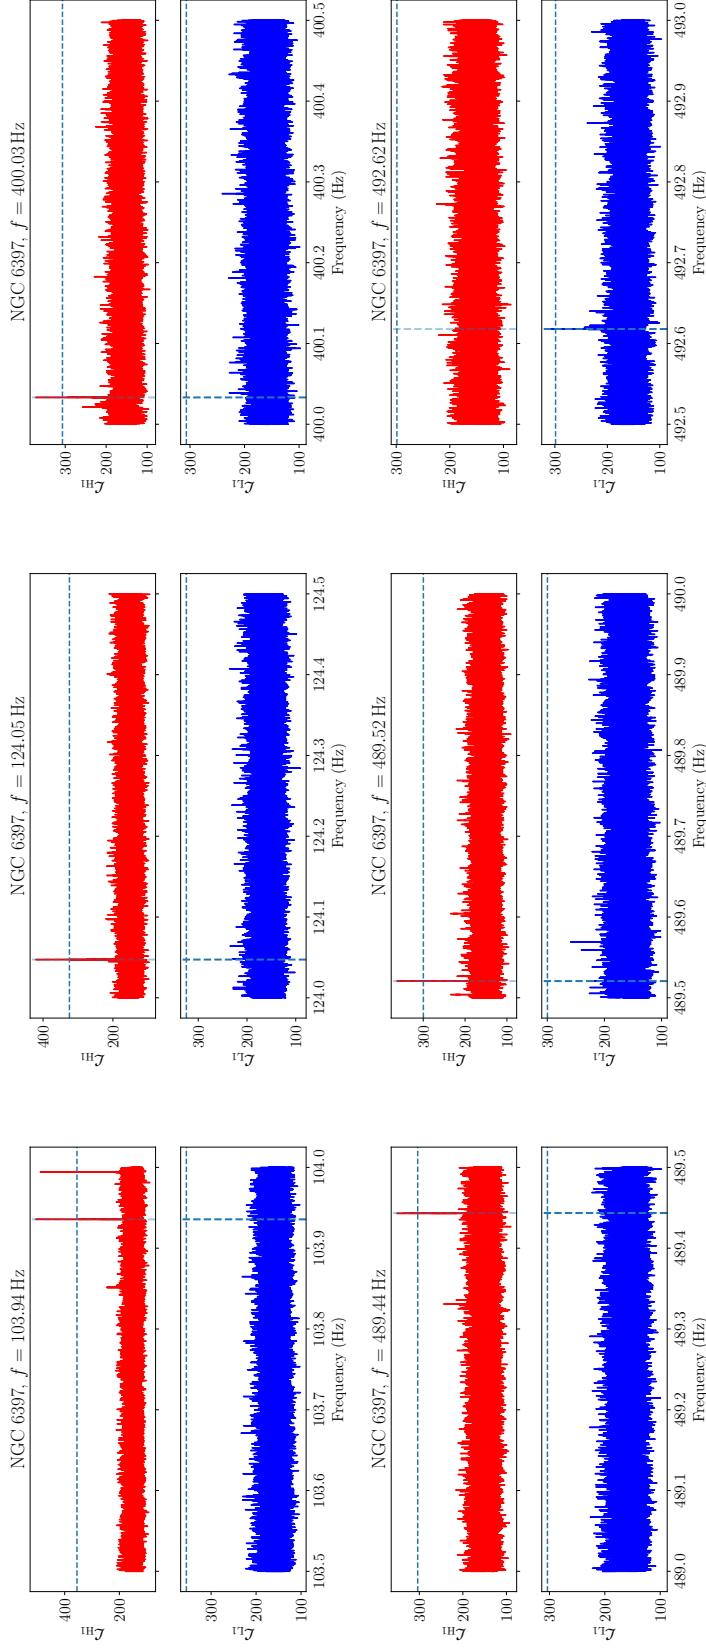


Figure 6.9: Single-detector log likelihoods \mathcal{L}_X versus frequency for the ten candidate groups surviving the cross-cluster veto. The vertical dashed lines indicate the frequency of the loudest candidate in each group. The horizontal dashed lines indicate the loudest joint-detector log likelihood in each group, \mathcal{L}_U . Nine of the ten candidate groups satisfy $\mathcal{L}_X > \mathcal{L}_U$ and are therefore vetoed. The cluster name and approximate edging frequency of the loudest candidate in each group are recorded above each panel.

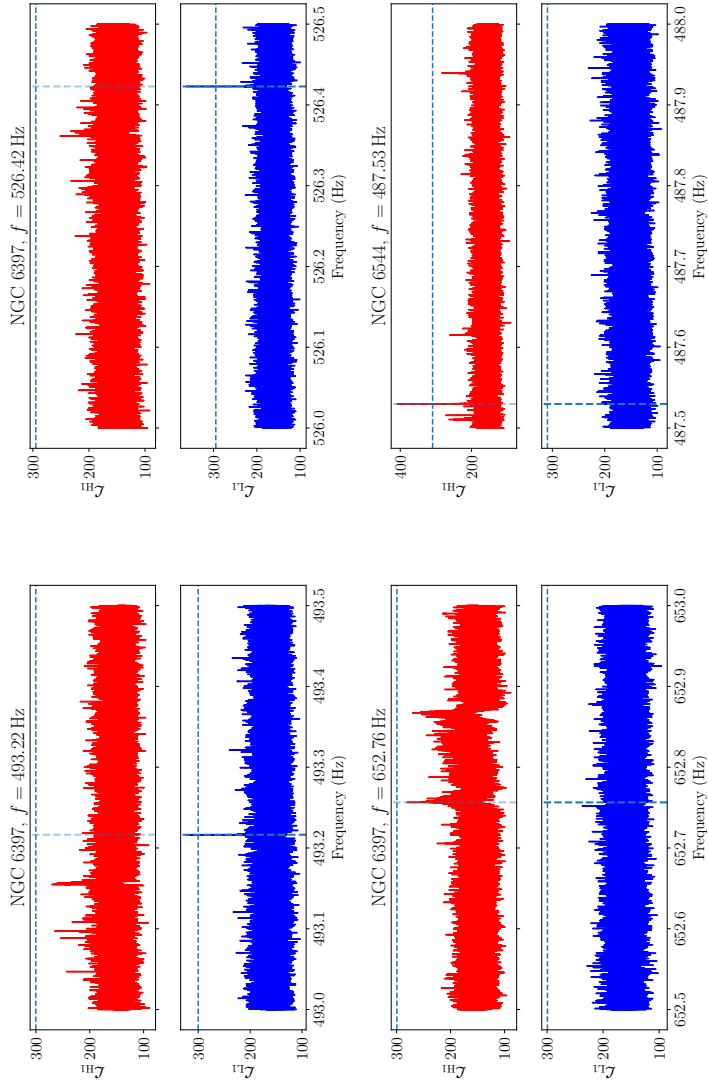


Figure 6.10: Single-detector log likelihoods \mathcal{L}_X versus frequency for the ten candidate groups surviving the cross-cluster veto. The vertical dashed lines indicate the frequency of the loudest candidate in each group. The horizontal dashed lines indicate the loudest joint-detector log likelihood in each group, \mathcal{L}_U . Nine of the ten candidate groups satisfy $\mathcal{L}_X > \mathcal{L}_U$ and are therefore vetoed. The cluster name and approximate edging frequency of the loudest candidate in each group are recorded above each panel. (Continued from Fig. 6.9).

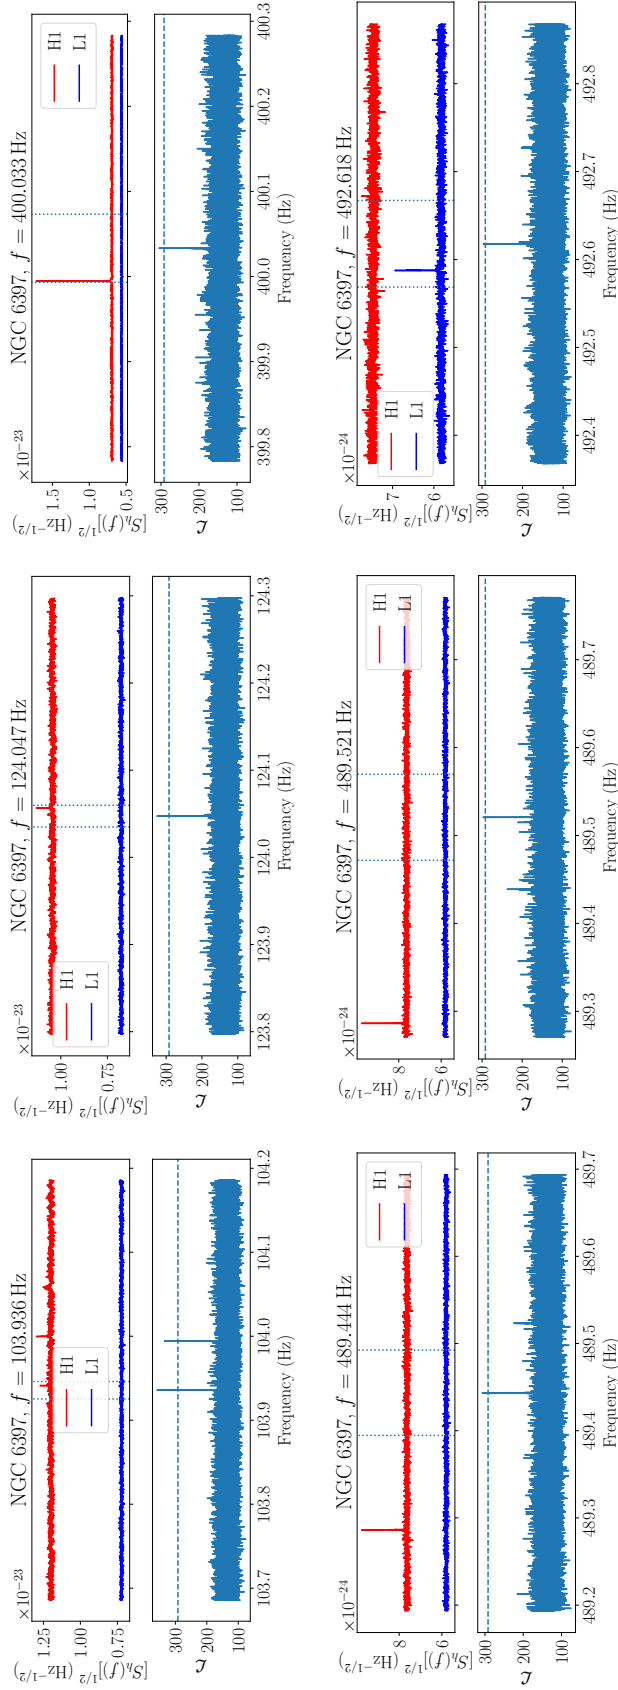


Figure 6.11: Single-detector ASDs (units: $\text{Hz}^{-1/2}$) and log likelihoods versus frequency (units: Hz) for the ten candidate groups surviving the cross-cluster veto. The vertical, dotted lines in the ASD plots indicate the extent of the Doppler modulation of the candidate group. The horizontal, dashed lines in the log likelihood plots indicate the value of \mathcal{L}_{th} for the candidate's cluster. Seven of the ten candidate groups show a clear disturbance in one of the detector ASDs. The cluster name and approximate ending frequency of the loudest path in each candidate group are recorded above each panel.




Figure 6.12: Single-detector ASDs (units: $\text{Hz}^{-1/2}$) and log likelihoods versus frequency (units: Hz) for the ten candidate groups surviving the cross-cluster veto. The vertical, dotted lines in the ASD plots indicate the extent of the Doppler modulation of the candidate group. The horizontal, dashed lines in the log likelihood plots indicate the value of \mathcal{L}_{th} for the candidate's cluster. Seven of the ten candidate groups show a clear disturbance in one of the detector ASDs. The cluster name and approximate ending frequency of the loudest path in each candidate group are recorded above each panel. (Continued from Fig. 6.11)

*Twenty miles left to the show
Hello, my old country, hello
Stars are just beginning to appear
And I have never in my life before been here*

Joanna Newsom

7

Conclusions and outlook

In this final chapter we summarise the results of Chapters 2–6 (Section 7.1), and look to possible directions for future work (Sections 7.2 and 7.3).

7.1 | Summary

The majority of the work in this thesis deals with the timing of radio pulsars, in particular the analysis of pulsar glitches using HMMs. Chapter 2 deals with issue of periodic observation scheduling and its interaction with glitch measurement (Section 7.1.1). Here the HMM glitch detector (Melatos et al., 2020) is used as a complementary method to traditional timing techniques, which is better able to alert the analyst to issues arising from periodic observation schedules. Chapter 3 presents a search for unknown pulsar glitches using the HMM glitch detector in 282 pulsars observed as part of the UTMOST-EW (UTMOST East-West) pulsar timing programme, and the computation of upper limits on the size of missed glitches in each pulsar (Section 7.1.2). Chapter 4 discusses a similar search for unknown glitches, but in a single high-quality dataset of observations of the Vela pulsar taken by the Mount Pleasant observatory. One new glitch is found, and upper limits are again set on the size of missed glitches (Section 7.1.3). Chapter 5 discusses initial science results from the UTMOST-NS (UTMOST North-South) pulsar timing programme using both the HMM glitch detector and ENTERPRISE (Ellis et al., 2020), including not only the glitches in the dataset but also the properties of the timing noise and the measured second frequency derivatives (Section 7.1.4). Finally, Chapter 6 discusses a different application of HMMs, to a search for continuous gravitational wave emission from neutron stars in nearby globular clusters (Section 7.1.5).

7.1.1 | Periodic observation scheduling and glitch measurement

In Chapter 2 we show that when pulsar timing observations are separated by multiples of a constant period, there is a degeneracy between pulsar timing models with signifi-

cantly different frequencies. We discuss this degeneracy in the context of the measurement of glitches. The frequency difference between degenerate models is approximately $1/T$, where T is the fundamental period of the observation schedule. We demonstrate through a case study of UTMOST observations of PSR J1709–4429 that this effect can lead to severe misestimation of glitch sizes, correcting the glitch size reported by Lower et al. (2020) of $\Delta\nu/\nu = (54.6 \pm 1.0) \times 10^{-9}$ to $\Delta\nu/\nu = (2432.2 \pm 0.1) \times 10^{-9}$. We also show that typical pulsar timing methods make it difficult to recognise this degeneracy. Though Bayesian pulsar timing methods could in principle alert the analyst by way of a multimodal posterior on the glitch size, in practice these methods rely on *pulse numbering* – assigning an absolute phase to each time of arrival *before* the inference step. This amounts to breaking the degeneracy between glitch models before the Bayesian analysis, and implicitly choosing a particular mode of the posterior. Usually this mode is the smallest positive frequency jump which is consistent with the data. The HMM method, on the other hand, does not require pulse numbering, and we show that this allows the posterior distribution on the frequency to reflect the degeneracy in the data, although the maximum *a posteriori* frequency estimates do not reliably correspond to the correct mode.

7.1.2 | Systematic glitch upper limits with an HMM

In Chapter 3 we demonstrate the utility of the HMM glitch detector as a means of understanding the completeness of glitch catalogues. This is essentially a question of interpreting non-detections. The HMM method provides a means of glitch detection which is deterministic, computationally inexpensive, and requires no input from the analyst except for initial configuration. This means that a question like “what is the probability of detecting a glitch of size $\Delta\nu = 10^{-8}$ Hz in a dataset that has the observing cadence and noise levels of the UTMOST-EW observations of PSR J1644–4559, using the HMM configured in a certain way?” has a meaningful answer, and it is feasible to estimate this answer empirically by injection studies. A more useful framing of this question is “at what glitch size would we expect to make a detection 90% of the time in this dataset, given an analyst-selected false alarm probability?” – this is the *upper limit* (at 90% confidence) on the size of any undetected glitches in the data.

Chapter 3 presents a search using the HMM glitch detector for previously unknown glitches and subsequent computation of upper limits on the size of undetected glitches in 282 datasets released as part of the first UTMOST data release (Lower et al., 2020). The upper limits vary between $4.1 \times 10^{-11} \leq \Delta\nu^{90\%}/\nu \leq 2.7 \times 10^{-7}$, with a mean value of 1.9×10^{-8} . We verify that including recovery has a modest effect on the upper limits – assuming complete recovery on a timescale of 100 d, the upper limits increase by a factor of $\lesssim 5$. Producing an upper limit on a single dataset involves analysing ~ 500 synthetic datasets with glitches injected; the full suite of upper limits thus amounts to an analysis of $\sim 1.4 \times 10^5$ datasets. The automated nature of the HMM glitch detector is crucial to enabling studies of this kind.

These upper limits constrain the possible glitching behaviour of the pulsars in question, but the observed glitch population includes a significant component which evades the upper limits derived here; see Figure 3.2. More constraining upper limits can be obtained through both higher quality datasets, in cases where the limiting factor is the ToA

uncertainty, and better understanding of the timing noise properties and incorporation of this understanding into the configuration of the HMM. Higher cadence observations do not necessarily improve the sensitivity to small glitches (see Section 4.5.2), but are invaluable in resolving short-timescale features in the post-glitch behaviour (e.g. quasi-exponential recovery).

7.1.3 | Searching for small glitches in the Vela pulsar

In Chapter 4 the emphasis shifts — rather than searching a large collection of datasets from many pulsars for glitches, we search for glitches in a single, high-quality dataset derived from dedicated monitoring of the Vela pulsar at the Mount Pleasant observatory covering 24 years from 1981 to 2005. However, the HMM still plays a central role in both the glitch search and setting of upper limits. We detect one previously unknown small glitch at MJD 48636, of fractional size $\Delta\nu/\nu = (8.19 \pm 0.04) \times 10^{-10}$, and with a fractional jump in frequency derivative $\Delta\dot{\nu}/\dot{\nu} = (2.98 \pm 0.1) \times 10^{-4}$. We also confirm previous detections of two other small glitches reported by Espinoza et al. (2021). The dataset is naturally broken into sections by the presence of large glitches and observing gaps. We analyse these sections separately and report per-section upper limits on the size of undetected glitches, with a median 90% upper limit on the fractional size of 1.35×10^{-9} .

The newly discovered glitch is the third smallest reported in Vela. Jankowski et al. (2015b) reported a glitch with $\Delta\nu/\nu = 4 \times 10^{-10}$, though subsequent analysis by Lower et al. (2020) has suggested that this may be a timing noise fluctuation. Zubieta et al. (2024) have recently reported an even smaller glitch, with $\Delta\nu/\nu = 2 \times 10^{-10}$. It is unclear at this stage at what point the timing noise in Vela overwhelms any glitch detection strategy — as discussed in Section 7.1.2, a more complete understanding of timing noise and its interaction with glitch detection methods is crucial to evaluating small glitches such as these, and to obtaining the most constraining upper limits on undetected glitches.

7.1.4 | Pulsar timing with UTMOST-NS

In Chapter 5 we widen our scope and present initial science results from the UTMOST-NS pulsar timing programme, which ran from 2021 to 2023. This includes not only an analysis of glitches in the dataset but also a discussion of the timing noise and braking indices in the sample. Where available the UTMOST-NS data are combined with earlier observations taken with UTMOST-EW, extending the timing baseline up to a maximum of 8 years.

We find that the contribution of timing noise to the residuals in each pulsar, denoted σ_{RN} , is correlated with the frequency derivative, confirming earlier studies (e.g. Shannon and Cordes, 2010; Lower et al., 2020) — assuming that σ_{RN} scales as $\chi_{\text{RN}} = \nu^a |\dot{\nu}|^b f_{\text{low}}^{-\gamma}$, we obtain $a = -0.92^{+0.31}_{-0.29}$, $b = 0.80^{+0.12}_{-0.13}$, and $\gamma = 1.8^{+0.6}_{-0.8}$, where f_{low} is the lowest-frequency mode included in the timing noise model. However, we also find that incorporating information about the spectral index of the timing noise power spectral density (PSD) does not improve the correlation between the degree of timing noise and the pulsar spin parameters — the estimated scatter in the relation between $\log_{10} \sigma_{\text{RN}}$ and $\log_{10} \chi_{\text{RN}}$ is $\epsilon = 0.69 \pm 0.09$ if γ is treated as a free parameter which is constant across the population, but increases to $\epsilon = 0.82 \pm 0.10$ if the value of γ is fixed to the theoretically expected

value of $(\beta - 1)/2$ for each pulsar, where β is the PSD spectral index of the timing noise. We also find that PSD parameters vary as more data are incorporated – with longer timing baselines there is a tendency to recover shallower PSDs.

Our investigation into second frequency derivatives confirms earlier studies by Vargas and Melatos (2023a) and Keith and Nițu (2023), both of which suggested that the reported measurements of $\ddot{\nu}$ from Bayesian pulsar timing packages such as TEMPONEST and ENTERPRISE must be regarded with some caution. In particular we demonstrate that the quoted uncertainties on $\ddot{\nu}$ can be underestimated by a factor of ≈ 3 if the pulsar has a relatively steep timing noise PSD, and a timing noise model is used which has a low-frequency cutoff at $1/T_{\text{span}}$ (which is the default behaviour in both TEMPONEST and ENTERPRISE). This underestimation of the $\ddot{\nu}$ uncertainty can lead to spurious “detections” of a non-zero $\ddot{\nu}$ – in our sample there are nine detections with implied braking indices ranging from $(-2.4 \pm 2.4) \times 10^4$ to $(5.0 \pm 0.7) \times 10^4$. When an extended timing noise model with a low-frequency cutoff at $1/(2T_{\text{span}})$ is included, all but one of these detections disappear, with the remaining case being PSR J0534+2200 – the Crab pulsar – for which we measure $n = 2.54 \pm 0.03$, in line with previous analyses (Lyne et al., 2015). We re-derive an analytic scaling between the magnitude of n and a combination of the timing noise strength and observing timespan (Vargas and Melatos, 2023a), and discuss the observing timespan required to measure the underlying secular braking index n_{pl} , assuming $1 \lesssim n_{\text{pl}} \lesssim 7$. For pulsars with $\beta \approx 4$ we find the median required observing timespan is 2.8×10^2 yr, while for pulsars with $\beta \approx 6$ we find a median required observing timespan of 3.9×10^7 yr.

The HMM glitch detector is again used to search for glitches in our sample, though there are some important differences to earlier applications. First, we use the HMM glitch detector in an “online” glitch detection pipeline, which ran on every new observation with a latency of ~ 1 hr, and produced an email alert in cases where a potentially significant glitch candidate was present. The three glitches which are known to have occurred during the operation of the UTMOST-NS timing programme were all detected by this pipeline (Dunn et al., 2021b; Dunn et al., 2022a; Dunn et al., 2023b). Second, the “offline” glitch search was run on the combined UTMOST-NS+EW dataset. This offline analysis is improved with respect to the search described in Chapter 3, primarily by the incorporation of information from ENTERPRISE about the timing noise characteristics, which allows us to tune the HMM parameters to accommodate the observed timing noise on a per-pulsar basis. The offline search does not detect any glitches which were not previously known from other observing campaigns or visual inspection of the timing residuals, and the mean upper limit in our sample is 5.3×10^{-9} (cf. 1.9×10^{-8} achieved for the search in Chapter 3), with a range $2.4 \times 10^{-12} \leq \Delta\nu^{90\%}/\nu \leq 1.3 \times 10^{-7}$. The seventeen known glitches in the sample are analysed using ENTERPRISE, and Bayesian model selection is used to determine whether an exponential recovery term is necessary. We detect recovery in seven glitches, with timescales between $1.9_{-0.6}^{+0.7}$ d and 613_{-353}^{+379} d. The three glitches in Vela all have recoveries with similar timescales, on the order of a few days – similar features have been noted in several previous Vela glitches (Flanagan, 1990; McCulloch et al., 1990; Dodson et al., 2002; Sarkissian et al., 2017)

7.1.5 | Continuous waves from neutron stars in globular clusters

Finally, we leave the domain of pulsar timing and look to gravitational waves as a potential source of insight on the nature of neutron stars. In Chapter 6 we report on a search for continuous gravitational waves from fast-spinning neutron stars in five nearby globular clusters. The search uses a HMM to track both the signal frequency and phase across coherent segments (Melatos et al., 2021), implementing the method for the first time in the context of an astrophysical search. We discuss the setup of the search in detail, including both the choice of various HMM parameters (e.g. number of phase bins and parameters specifying the transition matrix), the choice of ranges and template spacing for the sky position and frequency derivative, and the choice of threshold log likelihood which a candidate detection must exceed. Vetoes are applied to exclude candidates above threshold which are consistent with non-astrophysical, non-Gaussian noise features in the data. Candidates are vetoed if they come from frequency sub-bands which produce candidate lists larger than 100 MB on the grounds that these sub-bands are strongly contaminated by loud noise features, or if they overlap with noise features which have been identified in advance and associated with a definite non-astrophysical origin. These two vetoes serve to exclude a vast majority of the candidates. Candidates are further vetoed if they appear in multiple clusters, if they increase in significance when analysed with data from only one interferometer, or if they overlap with visually identifiable features in the noise amplitude spectral densities of each detector. After these veto procedures, no potential detections survive.

For each of the five clusters we estimate the sensitivity of the search via synthetic data injections, and convert these estimated strain sensitivities into a constraint on the ellipticity of neutron stars in the cluster. The best strain sensitivity is achieved for the cluster NGC 6544, with an upper limit of $h_{0,\text{eff}}^{95\%} = 2.7 \times 10^{-26}$ at 211 Hz assuming optimal orientation. The most constraining ellipticity limit is set for the cluster NGC 6397, with $\epsilon^{95\%} = 1.6 \times 10^{-7}$ at 800 Hz. These limits are a significant improvement over previous searches targeting globular clusters. The previous best 95% confidence upper limit on h_0 on CW emission from NGC 6544 is 6×10^{-25} at 170 Hz, set by Abbott et al. (2017d). Dergachev et al. (2019) performed a search targeting Terzan 5 (which is not targeted in the present search) and obtained a best upper limit on h_0 of 9×10^{-26} at 475 Hz.

7.2 | Future work: radio pulsars

The application of hidden Markov models to radio pulsar astronomy is still in relative infancy, and there are several directions for possible future developments in this area.

7.2.1 | Tuning the HMM glitch detector

In early applications of the HMM glitch detector the timing noise model was tuned for each pulsar to some extent according to rough scaling relations, but not directly adjusted based on any measured properties of the timing noise in each dataset (see for example Chapter 3). We have come to understand better how to make a connection between timing noise parameters phrased in terms of a PSD, and the timing noise model used in the HMM glitch detector — this is discussed in Chapter 5. However, this is not the

end of the issue. For instance: is ENTERPRISE (or an equivalent package) needed at all, or can we measure properties of the timing noise directly using the HMM pulsar timing framework? In a certain sense this appears straightforward – if the HMM is denoted by $\mathcal{M}(\phi)$ where ϕ parametrises the timing noise model, and the observed data are denoted by O , we can use the forward algorithm (Rabiner, 1989; Melatos et al., 2020) to calculate $\Pr[O | \mathcal{M}(\phi)]$, which serves as a likelihood function which we can use to obtain a posterior distribution on ϕ . But this approach has not been tested on either synthetic or real data, and if it does work in simple cases, it is unclear whether it can be extended to more complex cases where the vagaries of real data may interfere. Related to this is the question of what the space of timing noise models ought to encompass. Thus far we have only considered very simple cases, where the timing noise is driven by a random walk in either $\dot{\nu}$ or ν . However, a more complicated class of model may well be appropriate, involving for example a mixture of random walks (Cordes and Downs, 1985; D’Alessandro et al., 1995), interactions between coupled components of the star (Baym et al., 1969a; Meyers et al., 2021a; Antonelli et al., 2023; Vargas and Melatos, 2023a), or more complicated forms for the autocorrelation of the stochastic driving term (Melatos and Link, 2014).

Setting the timing noise model in the HMM by reference to a measured PSD presents its own issues – the currently implemented timing noise models give rise to very specific PSDs, which do not always coincide with those measured – in particular the spectral indices of the PSDs implied by the current HMM timing noise models take on only certain specific values, while measured spectral indices vary continuously. In Chapter 5 we sidestepped this issue by choosing the “closest” timing noise model for a given measured PSD, but this solution is somewhat *ad hoc*.

There is also the issue of small glitches and their effect on the measured timing noise properties. If we implement a timing noise model in the HMM by reference to a measured PSD, that PSD was necessarily obtained under the assumption that there are no small missed glitches in the dataset. If in fact there *are* small glitches present in the data which are not accounted for, how does this affect the PSD and hence the timing noise model incorporated into the HMM? Under what circumstances are we liable to miss small glitches due to this effect? If we aim to characterise the timing noise using the HMM alone, it is likely necessary to jointly search for small glitches and estimate the timing noise properties – is it feasible to do so?

Independent methods based not on HMMs but Kalman filtering may offer some paths forward as well (Meyers et al., 2021a; Meyers et al., 2021b; O’Neill et al., 2024). These techniques describe timing noise in terms of the stochastic evolution of the pulsar’s spin, which is readily translated to the HMM context, and have already been used to measure properties of the timing noise in individual pulsars (O’Neill et al., 2024). Similar techniques have also been applied to estimation of properties of systems containing accretion-powered pulsars (Melatos et al., 2023; O’Leary et al., 2024a; O’Leary et al., 2024b). While more work is needed to understand precisely how to translate the results from one method to the other, and the question of contamination by small glitches still remains, this is nonetheless a promising avenue for future applications.

7.2.2 | Discovering binary pulsars with an HMM

Another avenue for application of HMMs to radio pulsar astronomy is not in the delicate art of pulsar timing, but the discovery of new pulsars (O’Leary et al., [in prep.](#)). The detection of new pulsars is in many ways akin to the detection of continuous gravitational waves. Given a dedispersed intensity time series, a pulsar signal shows up in the Fourier spectrum as a series of peaks, with the fundamental harmonic at the observed pulsation frequency of the pulsar¹. For isolated pulsars with fast rotation frequencies the search process is then quite straightforward — one can form the Fourier spectrum of the dedispersed time series and look for significant peaks, since the frequency modulation due to intrinsic evolution of the pulsar is negligible over the timescales of an observation (\sim hours), and the Doppler modulation due to the motion of the Earth is easily removed in a pre-processing step.

The case of a binary pulsar, however, is more difficult. Almost always the binary parameters are *a priori* unknown, and the Doppler modulation due to the binary motion cannot be neglected — the natural frequency resolution for a 1 hour observation is $\sim 2 \times 10^{-4}$ Hz, and for a 100 Hz pulsar orbiting with a period of 1 d and a separation of 1 lt-s [which is well within the typically observed range (Manchester et al., [2005](#))] the magnitude of the Doppler shift is 7×10^{-3} Hz. Most searches for binary pulsars are therefore restricted to observations which cover only a small fraction, approximately 10%, of the pulsar’s binary orbit (Ransom, [2001](#)). With this restriction the acceleration of the pulsar is approximately constant, so that the frequency evolves linearly over the observation time, and a search over frequency and acceleration is feasible. But as observations grow longer, the curvature of the frequency track becomes important, and new techniques are needed (Andersen and Ransom, [2018](#); Balakrishnan et al., [2022](#); White et al., [2024](#)). One possible way forward is to construct a simple HMM which tracks the frequency evolution over the observation, in a spirit similar to the application to continuous gravitational wave searches — see the left panel of Figure 7.1. A simple transition matrix which allows for increments or decrements by one frequency bin each timestep can track the deterministic Doppler modulation, as long as the maximum change in frequency across one timestep does not exceed the frequency bin size. For pulsars in binary orbits with periods $\lesssim 1$ hr, the ability to track the binary modulation removes the constraint on observation timespan which applies to acceleration searches — e.g. for a pulsar with a 1 hr orbit, acceleration-based searches are unable to search more than approximately 10 minutes of data at a time without losing significant sensitivity, whereas an HMM-based approach can analyse a full pointing which might consist of hours of data. While for extremely short orbits other techniques have been developed (Ransom et al., [2003](#)), an HMM-based search could fill an important niche in the parameter space, where the orbital period within a factor of a few of the observation timespan. We note that the approach discussed here is quite similar to the Dynamic Power Spectrum method introduced by Chandler ([2003](#)), although this method has not been widely adopted.

¹Higher harmonics are present because the pulse profile is not a perfect sinusoid (Ransom et al., [2003](#); Yu, [2018](#)) — this effect is important, but we ignore it for this brief discussion.

7.2.3 | Discovering binary pulsars with an HMM: proof of concept

To end this section, we present a brief proof-of-concept demonstration of the above approach to searching for binary pulsars. The setup of the HMM is as follows (see Section 1.8.1 for a discussion of the HMM formalism):

- The hidden state of the system is the frequency of the pulsar as observed at the solar system barycenter.
- Given a dedispersed and barycentred time series generated from a telescope pointing, we divide the time series up into chunks of length T_{coh} , and compute the discrete Fourier transform for each chunk, retaining only the Fourier coefficients for frequencies within a user-specified range. The squared absolute values of these Fourier coefficients (i.e. the Fourier powers) serve as the emission (log) likelihoods.
- We assume a simple transition matrix where the frequency of the pulsar is allowed to jump by up to one bin in either direction at each timestep. This accommodates the deterministic sinusoidal binary Doppler modulation, as long as the maximum change in frequency between consecutive timesteps is no larger than a frequency bin ($1/T_{\text{coh}}$).
- We use the Viterbi algorithm to infer the hidden state evolution over the full observation.

Figure 7.1 shows a proof-of-concept example of this method of searching for new pulsars. We use `SIMULATESEARCH` (Hobbs et al., 2022) to simulate a 2-hour pointing with an injected signal from a 173 Hz pulsar in an 80-minute circular orbit with an orbital separation of 1.1 lt-s, corresponding to a companion mass of $\sim 5 M_{\odot}$ assuming a pulsar mass of $1.44 M_{\odot}$. We use `PRESTO` (Ransom, 2011) to process the simulated observation and generate the dedispersed and barycentred time series. We subsequently divide the time series into $T_{\text{coh}} = 30$ s chunks, and compute the Fourier powers for each chunk, keeping all frequency bins in the range 100–200 Hz. The frequency bin spacing is $1/(30$ s). The left panel shows the spectrogram for a loud injection, zoomed in on the location of the injected signal, along with the loudest Viterbi path (offset from the injection for visual clarity). The presence of the injected signal is clear from the spectrogram, and the Viterbi path follows the binary frequency modulation accurately, with a root mean square deviation of approximately 0.1 frequency bins from the true frequency path over the observation. The middle panel shows an equivalent spectrogram for a quieter injection, where the presence of a signal is not as clear from visual inspection of the spectrogram alone. The right panel again shows the spectrogram for the quiet injection, with the loudest Viterbi path shown in red (as before this is the loudest path across the full 100 Hz search band, although only a zoomed-in 5 Hz section is displayed in the heatmap), and the true frequency evolution shown as the dashed line. Although the path deviates from the injection somewhat towards the end of the observation, the signal is still successfully tracked, with a root mean square deviation of approximately 1 frequency bin from the true frequency path over the observation. The threshold signal-to-noise ratio above

which frequency tracking is still possible is of significant interest, but there is no simple way to calculate this threshold, and an empirical determination via synthetic signal injections is beyond the scope of the proof-of-concept work presented here.

7.3 | Future work: continuous gravitational waves

What next for hidden Markov models in the continuous wave context? Of course, new data will continue to come in – at the time of writing the second half of LIGO-Virgo-KAGRA’s fourth observing run is underway, and is expected to run until June 2025 (Shoemaker et al., 2024). Sensitivity improvements will come with each new observing run, and eventually the next generation of ground-based interferometers will come online and boost sensitivity substantially (Maggiore et al., 2020; Evans et al., 2023). The strain sensitivity frontier will thus be pushed downward without any significant new advances in how CW searches are performed. However, there is scope for input from the data analysis and astrophysics sides to accelerate this progress, and to provide direction to the limited resources available in order to maximize the chance of a first detection. In the rest of this section we discuss three potential paths for future work. Although they are framed by the search presented in Chapter 6, they are not specific to the details of this search.

7.3.1 | Follow-up analyses and the role of the phase-tracking HMM

In Chapter 6 we discussed a directed search for CWs from neutron stars in five globular clusters, using a phase-tracking HMM. The choice to use the phase-tracking HMM was motivated to a significant degree by a desire to see this new technique put to work in a full end-to-end search. In this respect the undertaking was successful – a new graphics processing unit-based implementation of the HMM and the phase-dependent \mathcal{B} -statistic was developed and tested as part of the project, and we have come to a better understanding of how to configure the search under certain assumptions about the potential spin wandering in the signal (see Section 6.3.4). However, it is not clear what the role of the phase-tracking HMM ought to be going forward. It is significantly more computationally expensive ($\sim 10\times$) on a per-template basis, and the sensitivity gains over the standard HMM implementation appear to be modest. For searches which are computationally limited, it may be that the phase-tracking HMM is best positioned as a followup tool rather than a first-line search method. In this case, we need to understand the behaviour of the phase-tracking HMM when it is focused on loud candidates from earlier search stages (whether that earlier stage is a standard HMM analysis or something else). What information do we actually gain by such a re-analysis, and how does it compare to other potential followup tools? There is no unique answer – followup methods are often tests for consistency of the candidate behaviour with the behaviour expected from an astrophysical signal², and hence necessarily involve assumptions about the signal model, which may differ depending on astrophysical context and the inclinations of the analyst. Nevertheless it is crucial to come to some kind of understanding of the proper use of the phase-tracking HMM to follow up signal candidates, and its value in this role. It is

²This is to be contrasted with veto procedures, which are tests for consistency with *non-Gaussian noise*.

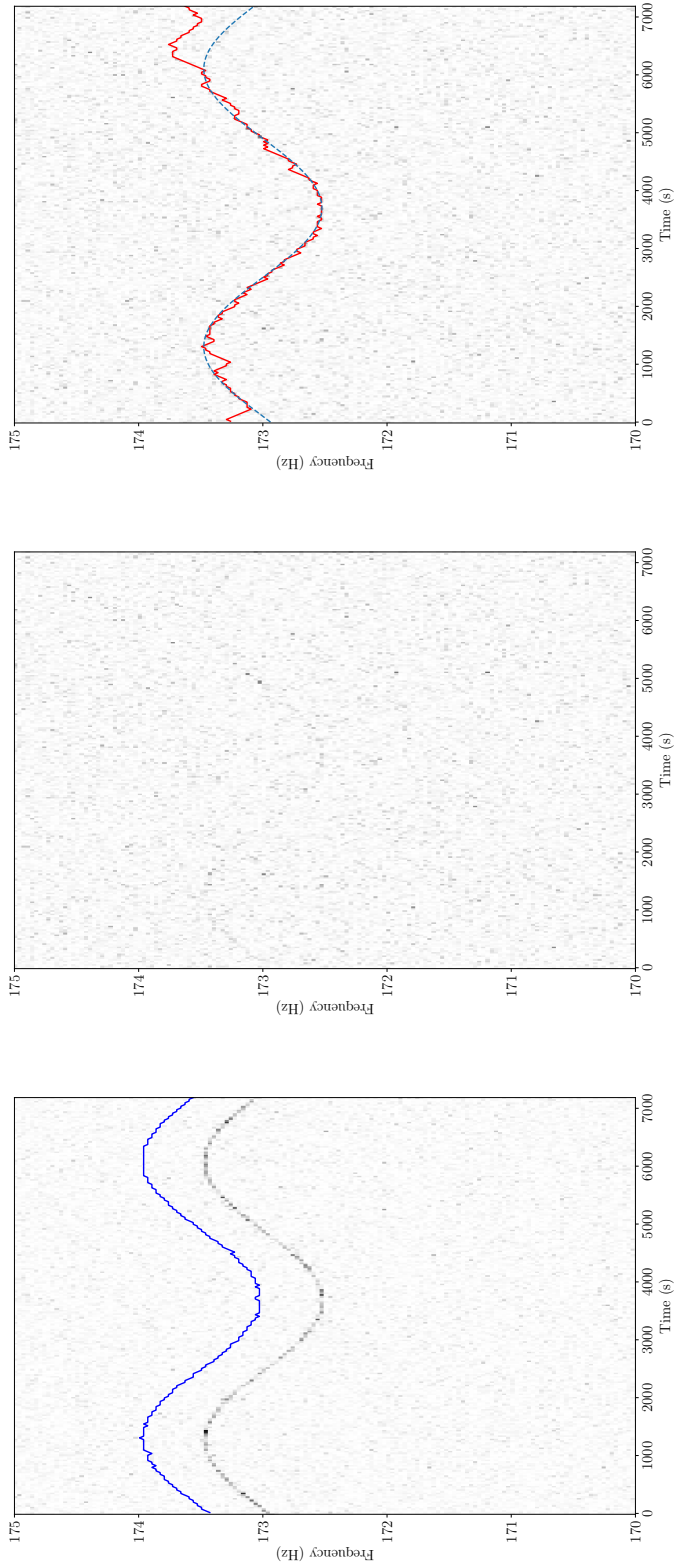


Figure 7.1: Demonstration of the use of HMMs for discovering binary pulsars, as discussed in Section 7.2.2. The three panels show spectrograms for a loud (*left*) and quiet (*middle and right*) injection of a 173 Hz pulsar in a 80-minute circular orbit. The injected signal is clear in the loud injection, and the blue track is the recovered Viterbi path (offset artificially by 0.5 Hz to make it visually distinguishable from the track in the spectrogram). The middle panel shows the spectrogram for the quiet injection with the same orbital parameters and with no Viterbi tracks overlaid, and the right panel shows the same spectrogram along with the recovered Viterbi path in red and the true path shown as the dashed line.

also valuable to consider how we might follow up candidates which are produced *by* the phase-tracking HMM. Appendix B briefly explores this question, presenting a possible follow-up method based on doubling the coherence time.

7.3.2 | Template banks for HMM searches

The search described in Chapter 6 involves a search over sky position (right ascension and declination) and spindown rate. As a simplification, we assess the “resolution” of the phase-tracking HMM in each of the three parameters by constructing a set of injected signals and scanning over the parameter of interest while holding the rest constant, empirically evaluating the loss of signal as a function of the difference between the parameter and the injected value. We then place templates in the desired parameter space volume as a simple rectilinear grid in the three parameters, with spacing informed by these resolution tests. This approach is suboptimal. There is a wealth of literature on the optimal construction of template banks for CW searches (see e.g. Prix and Itoh, 2005; Prix, 2007a; Prix, 2007b; Leaci and Prix, 2015; Wette, 2015; Wette et al., 2018), but the correct way to apply this work to the HMM context is unclear. The inclusion of stochastic spin wandering makes an analytic treatment of the “metric” structure of the parameter space challenging, and this impedes the construction of optimal template banks. Investigation into this issue is important not only for reducing the computational cost of searches. It also goes to the question of statistical correlation between nearby templates, which informs questions of detection significance (Wette, 2009). Even if an analytic approach is not feasible, more careful empirical analysis can still provide useful information about the parameter space structure and the optimal placement of templates.

7.3.3 | Globular clusters as CW search targets

The choice of targets in Chapter 6 raises two nested questions: why globular clusters? And why *those* globular clusters? A sharpening of our understanding of globular clusters as potential hosts of CW sources will inform the allocation of resources and choice of parameter ranges going forward. Globular clusters are of interest broadly speaking because they are small sky regions which are known to be rich in fast-spinning neutron stars, observed both as low-mass X-ray binaries (LMXBs) and radio millisecond pulsars (MSPs) (Clark, 1975; Ivanova et al., 2008; Ye et al., 2019). The case for CW emission from LMXBs is relatively clear, as they are actively accreting systems. Once accretion ceases, the lifetime and magnitude of any residual ellipticity is an area of active research (Vigelius and Melatos, 2009; Mukherjee et al., 2013; Suvorov and Melatos, 2019; Fujisawa et al., 2022) and these considerations may significantly impact the overall value of searches which target isolated neutron stars. We note, however, that MSPs may also experience transient accretion events after their LMXB phase, e.g. from asteroids or other debris whose orbit around the neutron star is disrupted (Cordes and Shannon, 2008; Shannon et al., 2013; Brook et al., 2014; Jennings et al., 2020). Synthesizing considerations of this kind into a coherent picture of the plausible CW source population in globular clusters will inform the extent to which it makes sense to allocate resources to searches to these targets, and what kinds of searches are best placed to make a detection – for example, does it make more sense to search for binary systems rather than isolated

neutron stars, in spite of the additional complexity incurred?

Taking for granted that globular clusters are broadly speaking a “good” target for CW searches, there is then the question of how to prioritise clusters in terms of detection probability. In Chapter 6 we took a simple approach which favoured clusters with high binary disruption rates, in order to produce a high fraction of isolated neutron stars post-recycling, since our search is not sensitive to sources which remain in binary systems, as the signal model does not include the Doppler modulation due to the binary orbit. But the clusters ranked highly by this metric do not necessarily host the most detectable continuous wave sources — there are additional factors involved, such as the overall production rate of MSPs in the cluster. A more complete investigation into the issue of cluster selection requires understanding not only how CW sources can be produced in globular clusters in the first place, but how observable features of those clusters (e.g. structural parameters such as core and half-light radii and core luminosity density) correlate to the processes which produce CW sources. This entails both observational studies (Verbunt and Hut, 1987; Boyles et al., 2011; Turk and Lorimer, 2013; Verbunt and Freire, 2014) and computational investigations of the properties of neutron star populations in globular clusters (Ivanova et al., 2008; Ye and Fragione, 2022).

*Oh me, oh my
I think it's been an eternity
You'd be surprised
At my degree of uncertainty*

Brian Eno



Bayesian pulsar timing

This appendix gives a brief overview of the pulsar timing likelihood which is implemented by Bayesian pulsar timing software such as `TEMPONEST` (Lentati et al., 2014) and `ENTERPRISE` (Ellis et al., 2020). It draws heavily on Johnson et al. (2024) and Taylor (2021), which discuss the approach to pulsar timing outlined here in more detail.

The N_{ToA} measured times of arrival, denoted \mathbf{t} , are decomposed into deterministic and stochastic contributions:

$$\mathbf{t} = \mathbf{t}^{\text{det}} + \mathbf{t}^{\text{sto}}, \quad (\text{A.1})$$

where \mathbf{t}^{det} is generated by a phase model like equation 1.11¹ (after transformation into the telescope frame) and \mathbf{t}^{sto} is the stochastic part generated by a combination of timing noise and instrumental measurement noise.

The stochastic contribution has a covariance matrix \mathbf{C} with elements

$$C_{ij} = \langle \mathbf{t}_i^{\text{sto}} \mathbf{t}_j^{\text{sto}} \rangle, \quad (\text{A.2})$$

where i, j label the observed ToAs and the likelihood of a given set of observations is given by

$$\Pr(\mathbf{t} | \boldsymbol{\theta}) = \frac{1}{\sqrt{2\pi \det \mathbf{C}}} \exp\left[-\frac{1}{2} (\mathbf{t} - \mathbf{t}^{\text{det}})^\top \mathbf{C}^{-1} (\mathbf{t} - \mathbf{t}^{\text{det}})\right], \quad (\text{A.3})$$

where $\boldsymbol{\theta}$ is a vector of parameters specifying \mathbf{t}^{det} and \mathbf{C} . Although this formulation is in principle complete, in practice it is unwieldy – in the presence of timing noise \mathbf{C} is dense, and evaluating the likelihood involves inverting this matrix, which is computationally expensive. It is therefore typical to approximate the contribution of the timing noise to \mathbf{t}^{sto} as a sum of sinusoids (Lentati et al., 2013), viz.

$$\mathbf{t}^{\text{sto}} = \mathbf{F}\mathbf{a} + \mathbf{n}, \quad (\text{A.4})$$

¹In many cases, particularly applications to pulsar timing array analyses, t_{det} is further decomposed into a form which is linearised about a fiducial timing model, but in the applications in this thesis we do not perform this decomposition.

where $\mathbf{F}\mathbf{a}$ is given by

$$(\mathbf{F}\mathbf{a})_i = \sum_{j=1}^N [X_j \sin(2\pi f_j t_i) + Y_j \cos(2\pi f_j t_i)], \quad (\text{A.5})$$

where the Fourier amplitudes \mathbf{X} and \mathbf{Y} are contained in the $2N \times 1$ vector \mathbf{a} and the $n \times 2N$ matrix \mathbf{F} contains the Fourier components evaluated at each of the ToAs. The values of f_j are chosen in advance and typically taken to be $f_j = j/T$, $j = 1, 2, \dots, N$ where T is the total length of the dataset and $N = O(10)$. The \mathbf{n} term contains the white noise contribution – for the purposes of this thesis it suffices to consider a diagonal² covariance matrix,

$$\mathbf{N} = \langle \mathbf{n}\mathbf{n}^\top \rangle = \text{diag}(\sigma_1^2, \sigma_2^2, \dots, \sigma_n^2) \quad (\text{A.6})$$

where σ_i is the uncertainty of the i th ToA. The likelihood then takes on the simpler form

$$\Pr(\mathbf{t} \mid \boldsymbol{\theta}^{\text{det}}, \mathbf{a}) = \frac{1}{\sqrt{2\pi \det \mathbf{N}}} \exp\left[-\frac{1}{2}(\mathbf{t} - \mathbf{t}^{\text{det}} - \mathbf{F}\mathbf{a})^\top \mathbf{N}^{-1}(\mathbf{t} - \mathbf{t}^{\text{det}} - \mathbf{F}\mathbf{a})\right] \quad (\text{A.7})$$

where now the only matrix inversion required is for the diagonal matrix \mathbf{N} .

We typically do not wish to estimate the Fourier amplitudes \mathbf{a} directly. Instead we make some assumption about the functional form of the power spectral density (PSD) $S(f)$, which is parametrised by a set of hyperparameters $\boldsymbol{\eta}$ which defines a prior distribution on \mathbf{a} , taken to be a multivariate Gaussian with covariance $\boldsymbol{\varphi}(\boldsymbol{\eta})$. Stationarity implies that the Fourier components are independent and we thus write

$$\boldsymbol{\varphi}_{jk} = \kappa_j \delta_{jk}, \quad (\text{A.8})$$

where δ_{jk} is the Kronecker delta and κ_j is determined by the PSD. In general we have $\kappa_j = S(f)\Delta f$ where typically $\Delta f = 1/T$. The most common parametrisation of the PSD is as a power law:

$$S(f) = (1 \text{ yr}^{-1})^3 \frac{A_{\text{red}}^2}{12\pi^2} \left(\frac{f}{1 \text{ yr}^{-1}}\right)^\beta, \quad (\text{A.9})$$

where the hyperparameters are $\boldsymbol{\eta} = (A_{\text{red}}, \beta)$. It is then possible to marginalise over the Fourier coefficients \mathbf{a} and obtain $\Pr(\mathbf{t} \mid \boldsymbol{\theta}^{\text{det}}, \boldsymbol{\eta})$, so that we can easily sample the timing noise hyperparameters directly.

²In some applications \mathbf{N} is only block-diagonal, but this still simplifies the problem greatly.

We didn't ask what it seems like, we asked what it is!

Aaron Weiss

B

A signal consistency test for continuous wave signals

B.1 | Vetoes and signal consistency tests

When faced with a candidate CW detection, a variety of follow-up procedures are available which aim to classify the candidate as either a noise fluctuation which is of no astrophysical significance, or a *bona fide* astrophysical signal. Many of these procedures are *veto*s – these are tests which check whether a candidate has characteristics which are consistent with a non-Gaussian noise artifact. Typically, vetoes are applied sequentially, and candidates which are flagged at any stage as being consistent with a non-Gaussian feature are removed from consideration from any subsequent analysis. Notice that there is nothing in the fundamental logic of the veto which refers to the qualities expected of an astrophysical signal. There is the question of *veto safety*, which asks how often a veto procedure will reject an astrophysical candidate (see e.g. Jones et al., 2022; Jaume et al., 2024) – but this is logically separate to the statement we make when applying a veto, namely “this candidate is consistent with a non-Gaussian noise feature” *not* “this candidate is inconsistent with an astrophysical signal”.

There is a second class of follow-up procedures which address the question left unanswered by vetoes: does the candidate behave like an astrophysical signal? We refer to procedures which aim to address this question as *signal consistency tests*. In order to specify such a test, we first need to decide on what counts as an astrophysical signal. There are myriad answers to this question, none of them uniquely correct. The analyst is obliged to make a choice here about the range of signal models their search will cover, which encompasses issues of both astrophysical prior belief and the scope of a given work. Nevertheless such choices can be made, and to some extent are implicitly made when the parameters of a search are set. Note, however, that signal consistency checks are not completely beholden to the signal model implied by the setup of the search. For example, in a standard Viterbi search the range of “astrophysical signals” which one

might include in a consistency check ought not be restricted to signals which change in frequency stepwise on the boundaries between the coherent steps.

B.2 | Doubling coherence time as a signal consistency test

In the remainder of this appendix we present one particular possible signal consistency check, in the context of the search for CW emission from globular clusters presented in Chapter 6. In an earlier iteration of that analysis¹, six candidates were found which could not be rejected on the basis of the veto procedures in place. The signal consistency test presented here was devised to determine whether these candidates fell within a class of plausible astrophysical signals – if they did not, then they were rejected.

The consistency test is based on the re-analysis of the candidate with a doubled coherence time T_{coh} . In the case of a semicoherent search and a signal with no spin wandering or other perturbations to the signal phase not captured by the model (e.g. long-period binary orbital modulation), we expect that the significance (suitably defined) of an astrophysical candidate will increase as the coherence time increases – this is the idea behind e.g. Markov Chain Monte Carlo techniques for following up CW candidates using the \mathcal{F} -statistic (Ashton and Prix, 2018; Tenorio et al., 2021) as well as a signal consistency test used in a Viterbi-based analysis of outliers from the O3 all-sky all-frequency radiometer search (Abbott et al., 2022d; Knee et al., 2024). However, the addition of spin wandering complicates matters – if the characteristics of the signal change too wildly, on too short a timescale, then increasing the coherence time will not improve the significance of the candidate (Carlin and Melatos, [submitted](#)). In the context of the search presented in Chapter 6, we do not wish to exclude spin wandering signals entirely – although the choice of coherence time in this search is primarily driven by the desire to keep the search computationally manageable (see Section 6.3), we nevertheless wish to maintain sensitivity to a broad range of possible signals, including those with a significant degree of spin wandering. When performing signal injections to assess the sensitivity of this search, we adopt a prescription where spin wandering is injected via a randomly chosen \ddot{f} , which changes stepwise every 7 days within a range

$$\pm\Delta\ddot{f} = \pm\Delta f / (4T_{\text{coh}}^2) = \pm 1.55 \times 10^{-18} \text{ Hz s}^{-2}, \quad (\text{B.1})$$

with $T_{\text{coh}} = 5 \text{ d}$ and $\Delta f = 1/2T_{\text{coh}}$.

How do candidates arising from this class of signal model behave when re-analysed after doubling T_{coh} ? Here we define the significance of a candidate via its “Viterbi score” (Abbott et al., 2017e), defined as

$$\mathcal{S} = \frac{\mathcal{L} - \mu_{\mathcal{L}}}{\sigma_{\mathcal{L}}}, \quad (\text{B.2})$$

where \mathcal{L} is the log likelihood of the candidate, $\mu_{\mathcal{L}}$ is the mean of the log likelihood distribution in noise, and $\sigma_{\mathcal{L}}$ is the standard deviation of the noise-only distribution. We then perform a suite of injections into the LIGO O3 data with spin wandering injected

¹Superseded due to bug fixes and accompanying changes to the template spacing. Note also that this earlier iteration was performed with a lower log likelihood threshold, corresponding to one false alarm per 100 Hz per cluster, instead of the one false alarm per cluster adopted in the final version of the search.

according to equation (B.1) and analyse each injection using both a 5 d and 10 d coherence time, searching over a fine grid in α , δ , and \dot{f} and using the loudest candidate in each case to determine \mathcal{S} . We choose random sky positions for each injection, and keep the initial frequency fixed at 321.3563573 Hz (chosen so that the generated data are not contaminated by obvious lines or other non-Gaussian features). The effective h_0 of the injected signal (see equation 6.25) is fixed at 2.3×10^{-26} , and the inclination angle is allowed to vary so that $\cos \iota$ is uniform on $[-1, 1]$.

For each injection we thus obtain a value of \mathcal{S} for both $T_{\text{coh}} = 5$ d and $T_{\text{coh}} = 10$ d. The top panel of Figure B.1 shows the results of 200 such injections, with $\mathcal{S}_{5\text{d}}$ on the horizontal axis and $\mathcal{S}_{10\text{d}}$ on the vertical axis, and the condition $\mathcal{S}_{5\text{d}} = \mathcal{S}_{10\text{d}}$ indicated by the dashed line. In all realisations we find that $\mathcal{S}_{10\text{d}} > \mathcal{S}_{5\text{d}}$. This suggests that $\mathcal{S}_{10\text{d}} > \mathcal{S}_{5\text{d}}$ is a useful condition for assessing whether a given candidate falls within this class of astrophysical signal models. The middle and bottom panels of Figure B.1 show the equivalent plots but in the case where $\Delta\dot{f}$ is halved and doubled, respectively. As expected, halving $\Delta\dot{f}$ increases the separation between $\mathcal{S}_{5\text{d}}$ and $\mathcal{S}_{10\text{d}}$, while doubling $\Delta\dot{f}$ brings them closer together – in a few cases $\mathcal{S}_{10\text{d}} > \mathcal{S}_{5\text{d}}$ is no longer satisfied. This gives us a rough understanding of what kinds of signals we might reject by enforcing the condition $\mathcal{S}_{10\text{d}} > \mathcal{S}_{5\text{d}}$.

B.3 | Application to candidates from superseded GC search

We provide for completeness the details of the six candidates which arose in the earlier iteration (see footnote 1) of the globular cluster search of Chapter 6, and motivated this investigation. Table B.1 lists the details of the candidates. Figures B.2–B.7 show the results of fine scans in α , δ , and \dot{f} around each candidate for $T_{\text{coh}} = 5$ d and $T_{\text{coh}} = 10$ d. The maximum Viterbi scores recorded are shown in the titles of the subpanels – in all cases the maximum $\mathcal{S}_{5\text{d}}$ exceeds the maximum $\mathcal{S}_{10\text{d}}$, violating the criterion established in this appendix $\mathcal{S}_{10\text{d}} > \mathcal{S}_{5\text{d}}$. As such, we do not regard these candidates as corresponding to astrophysical signals with spin wandering at or below the level of $\Delta\dot{f} = \pm 1.55 \times 10^{-18} \text{ Hz s}^{-2}$, and do not consider them further.




Figure B.1: Comparison of Viterbi scores \mathcal{S} obtained for $T_{\text{coh}} = 5$ d and 10 d in a set of 200 injections with \dot{f} chosen at random every 7 days in a uniform range centred on zero with extent $\pm 1.55 \times 10^{-18} \text{ Hz s}^{-2}$ (*top*), $\pm 0.76 \times 10^{-18} \text{ Hz s}^{-2}$ (*middle*), and $\pm 3.1 \times 10^{-18} \text{ Hz s}^{-2}$ (*bottom*). The dashed line indicates $\mathcal{S}_{5\text{d}} = \mathcal{S}_{10\text{d}}$.

Table B.1: Details of the six candidates arising from an earlier iteration of the search for CW emission from globular clusters detailed in Chapter 6. For each candidate we give the corresponding cluster, the right ascension and declination (J2000), the ending frequency of the loudest Viterbi path, the value of f for the loudest candidate, and the log likelihood \mathcal{L} .

| Cluster | Right ascension | Declination | Ending f (Hz) | f (Hz s ⁻¹) | \mathcal{L} |
|----------|---|-------------|-------------------|-----------------------------|---------------|
| NGC 6540 | 18 ^h 6 ^m 49.5 ^s | -27°45'55" | 460.2541164351852 | -4.20 × 10 ⁻¹⁰ | 293.7 |
| NGC 6325 | 17 ^h 17 ^m 58.7 ^s | -23°45'58" | 293.6517611111111 | -2.805 × 10 ⁻¹⁰ | 267.8 |
| NGC 6325 | 17 ^h 17 ^m 21.5 ^s | -23°45'58" | 320.692593287037 | -1.1945 × 10 ⁻¹⁰ | 268.0 |
| NGC 6397 | 17 ^h 40 ^m 20.1 ^s | -53°51'22" | 491.2474289351852 | -1.55 × 10 ⁻¹⁰ | 285.8 |
| NGC 6544 | 18 ^h 7 ^m 20.6 ^s | -24°59'50" | 588.0081222222223 | -2.4575 × 10 ⁻¹⁰ | 265.2 |
| Terzan 6 | 17 ^h 50 ^m 12.2 ^s | -31°30'20" | 702.7692078703705 | -2.6496 × 10 ⁻¹⁰ | 285.7 |




Figure B.2: Heatmaps of log likelihood \mathcal{L} for a fine-grid search in α , δ , and \dot{f} around the candidate signal in Terzan 6 at $f \approx 702$ Hz found in an early iteration of the globular cluster CW search. The top two panels present the results for $T_{\text{coh}} = 5$ d, while the bottom two present the results for $T_{\text{coh}} = 10$ d. The starred quantities \dot{f}^* , α^* , and δ^* are the values corresponding to the loudest candidate in each search. The maximum Viterbi scores in each case are given in the titles of the subfigures. The candidate is rejected because the 10 d Viterbi scores do not exceed the 5 d scores — the features of the heatmaps are shown for illustrative purposes but are not the basis for any assessment of the candidate.






Figure B.3: As in Figure B.2 but for the candidate in NGC 6325 at $f \approx 320$ Hz.


NGC 6325, $f = 293.652$ Hz, $T_{\text{coh}} = 5$ d, Viterbi score: 10.07



NGC 6325, $f = 293.652$ Hz, $T_{\text{coh}} = 5$ d, Viterbi score: 10.07



NGC 6325, $f = 293.652$ Hz, $T_{\text{coh}} = 10$ d, Viterbi score: 9.25



NGC 6325, $f = 293.652$ Hz, $T_{\text{coh}} = 10$ d, Viterbi score: 9.25






Figure B.4: As in Figure B.2 but for the candidate in NGC 6325 at $f \approx 293$ Hz.


NGC 6544, $f = 588.008$ Hz, $T_{\text{coh}} = 5$ d, Viterbi score: 8.47



NGC 6544, $f = 588.008$ Hz, $T_{\text{coh}} = 5$ d, Viterbi score: 8.47



NGC 6544, $f = 588.008$ Hz, $T_{\text{coh}} = 10$ d, Viterbi score: 8.02



NGC 6544, $f = 588.008$ Hz, $T_{\text{coh}} = 10$ d, Viterbi score: 8.02






Figure B.5: As in Figure B.2 but for the candidate in NGC 6544 at $f \approx 588$ Hz.


NGC 6397, $f = 491.247$ Hz, $T_{\text{coh}} = 5$ d, Viterbi score: 10.67



NGC 6397, $f = 491.247$ Hz, $T_{\text{coh}} = 5$ d, Viterbi score: 10.67



NGC 6397, $f = 491.247$ Hz, $T_{\text{coh}} = 10$ d, Viterbi score: 10.58



NGC 6397, $f = 491.247$ Hz, $T_{\text{coh}} = 10$ d, Viterbi score: 10.58






Figure B.6: As in Figure B.2 but for the candidate in NGC 6397 at $f \approx 491$ Hz.


NGC 6540, $f = 460.254$ Hz, $T_{\text{coh}} = 5$ d, Viterbi score: 11.48



NGC 6540, $f = 460.254$ Hz, $T_{\text{coh}} = 5$ d, Viterbi score: 11.48



NGC 6540, $f = 460.254$ Hz, $T_{\text{coh}} = 10$ d, Viterbi score: 9.46



NGC 6540, $f = 460.254$ Hz, $T_{\text{coh}} = 10$ d, Viterbi score: 9.46




Figure B.7: As in Figure B.2 but for the candidate in NGC 6540 at $f \approx 460$ Hz.

Bibliography

- Aasi, J. et al. (Nov. 2013). “Directed search for continuous gravitational waves from the Galactic center”. In: *Physical Review D* 88.10, 102002, p. 102002. DOI: [10.1103/PhysRevD.88.102002](https://doi.org/10.1103/PhysRevD.88.102002). arXiv: [1309.6221](https://arxiv.org/abs/1309.6221) [gr-qc].
- (Apr. 2014). “Gravitational Waves from Known Pulsars: Results from the Initial Detector Era”. In: *The Astrophysical Journal* 785.2, 119, p. 119. DOI: [10.1088/0004-637X/785/2/119](https://doi.org/10.1088/0004-637X/785/2/119). arXiv: [1309.4027](https://arxiv.org/abs/1309.4027) [astro-ph.HE].
- Aasi, J. et al. (Feb. 2016). “Search of the Orion spur for continuous gravitational waves using a loosely coherent algorithm on data from LIGO interferometers”. In: *Physical Review D* 93.4, 042006, p. 042006. DOI: [10.1103/PhysRevD.93.042006](https://doi.org/10.1103/PhysRevD.93.042006). arXiv: [1510.03474](https://arxiv.org/abs/1510.03474) [gr-qc].
- Abbott, B. et al. (Jan. 2008a). “All-sky search for periodic gravitational waves in LIGO S4 data”. In: *Physical Review D* 77.2, 022001, p. 022001. DOI: [10.1103/PhysRevD.77.022001](https://doi.org/10.1103/PhysRevD.77.022001). arXiv: [0708.3818](https://arxiv.org/abs/0708.3818) [gr-qc].
- Abbott, B. et al. (Aug. 2008b). “Beating the Spin-Down Limit on Gravitational Wave Emission from the Crab Pulsar”. In: *The Astrophysical Journal* 683.1, p. L45. DOI: [10.1086/591526](https://doi.org/10.1086/591526). arXiv: [0805.4758](https://arxiv.org/abs/0805.4758) [astro-ph].
- Abbott, B. P. et al. (Apr. 2010). “Searches for Gravitational Waves from Known Pulsars with Science Run 5 LIGO Data”. In: *The Astrophysical Journal* 713.1, pp. 671–685. DOI: [10.1088/0004-637X/713/1/671](https://doi.org/10.1088/0004-637X/713/1/671). arXiv: [0909.3583](https://arxiv.org/abs/0909.3583) [astro-ph.HE].
- Abbott, B. P. et al. (Apr. 2016a). “GW150914: The Advanced LIGO Detectors in the Era of First Discoveries”. In: *Physical Review Letters* 116.13, 131103, p. 131103. DOI: [10.1103/PhysRevLett.116.131103](https://doi.org/10.1103/PhysRevLett.116.131103). arXiv: [1602.03838](https://arxiv.org/abs/1602.03838) [gr-qc].
- (Feb. 2016b). “Observation of Gravitational Waves from a Binary Black Hole Merger”. In: *Physical Review Letters* 116.6, 061102, p. 061102. DOI: [10.1103/PhysRevLett.116.061102](https://doi.org/10.1103/PhysRevLett.116.061102). arXiv: [1602.03837](https://arxiv.org/abs/1602.03837) [gr-qc].
- (Apr. 2017a). “First Search for Gravitational Waves from Known Pulsars with Advanced LIGO”. In: *The Astrophysical Journal* 839.1, 12, p. 12. DOI: [10.3847/1538-4357/aa677f](https://doi.org/10.3847/1538-4357/aa677f). arXiv: [1701.07709](https://arxiv.org/abs/1701.07709) [astro-ph.HE].
- Abbott, B. P. et al. (Oct. 2017b). “GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral”. In: *Physical Review Letters* 119.16, 161101, p. 161101. DOI: [10.1103/PhysRevLett.119.161101](https://doi.org/10.1103/PhysRevLett.119.161101). arXiv: [1710.05832](https://arxiv.org/abs/1710.05832) [gr-qc].
- (Oct. 2017c). “Multi-messenger Observations of a Binary Neutron Star Merger”. In: *The Astrophysical Journal* 848.2, L12, p. L12. DOI: [10.3847/2041-8213/aa91c9](https://doi.org/10.3847/2041-8213/aa91c9). arXiv: [1710.05833](https://arxiv.org/abs/1710.05833) [astro-ph.HE].
- Abbott, B. P. et al. (Apr. 2017d). “Search for continuous gravitational waves from neutron stars in globular cluster NGC 6544”. In: *Physical Review D* 95.8, 082005, p. 082005. DOI: [10.1103/PhysRevD.95.082005](https://doi.org/10.1103/PhysRevD.95.082005). arXiv: [1607.02216](https://arxiv.org/abs/1607.02216) [gr-qc].

- Abbott, B. P. et al. (June 2017e). “Search for gravitational waves from Scorpius X-1 in the first Advanced LIGO observing run with a hidden Markov model”. In: *Physical Review D* 95.12, 122003, p. 122003. DOI: [10.1103/PhysRevD.95.122003](https://doi.org/10.1103/PhysRevD.95.122003). arXiv: [1704.03719](https://arxiv.org/abs/1704.03719) [gr-qc].
- Abbott, B. P. et al. (July 2019a). “GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs”. In: *Physical Review X* 9.3, 031040, p. 031040. DOI: [10.1103/PhysRevX.9.031040](https://doi.org/10.1103/PhysRevX.9.031040). arXiv: [1811.12907](https://arxiv.org/abs/1811.12907) [astro-ph.HE].
- (June 2019b). “Narrow-band search for gravitational waves from known pulsars using the second LIGO observing run”. In: *Physical Review D* 99.12, 122002, p. 122002. DOI: [10.1103/PhysRevD.99.122002](https://doi.org/10.1103/PhysRevD.99.122002). arXiv: [1902.08442](https://arxiv.org/abs/1902.08442) [gr-qc].
 - (Dec. 2019c). “Search for gravitational waves from Scorpius X-1 in the second Advanced LIGO observing run with an improved hidden Markov model”. In: *Physical Review D* 100.12, 122002, p. 122002. DOI: [10.1103/PhysRevD.100.122002](https://doi.org/10.1103/PhysRevD.100.122002). arXiv: [1906.12040](https://arxiv.org/abs/1906.12040) [gr-qc].
 - (July 2019d). “Searches for Gravitational Waves from Known Pulsars at Two Harmonics in 2015-2017 LIGO Data”. In: *The Astrophysical Journal* 879.1, 10, p. 10. DOI: [10.3847/1538-4357/ab20cb](https://doi.org/10.3847/1538-4357/ab20cb). arXiv: [1902.08507](https://arxiv.org/abs/1902.08507) [astro-ph.HE].
- Abbott, R. et al. (Oct. 2021a). “All-sky search for continuous gravitational waves from isolated neutron stars in the early O3 LIGO data”. In: *Physical Review D* 104.8, 082004, p. 082004. DOI: [10.1103/PhysRevD.104.082004](https://doi.org/10.1103/PhysRevD.104.082004). arXiv: [2107.00600](https://arxiv.org/abs/2107.00600) [gr-qc].
- (Mar. 2021b). “All-sky search in early O3 LIGO data for continuous gravitational-wave signals from unknown neutron stars in binary systems”. In: *Physical Review D* 103.6, 064017, p. 064017. DOI: [10.1103/PhysRevD.103.064017](https://doi.org/10.1103/PhysRevD.103.064017). arXiv: [2012.12128](https://arxiv.org/abs/2012.12128) [gr-qc].
 - (Apr. 2021c). “GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo during the First Half of the Third Observing Run”. In: *Physical Review X* 11.2, 021053, p. 021053. DOI: [10.1103/PhysRevX.11.021053](https://doi.org/10.1103/PhysRevX.11.021053). arXiv: [2010.14527](https://arxiv.org/abs/2010.14527) [gr-qc].
 - (Nov. 2021d). “Searches for Continuous Gravitational Waves from Young Supernova Remnants in the Early Third Observing Run of Advanced LIGO and Virgo”. In: *The Astrophysical Journal* 921.1, 80, p. 80. DOI: [10.3847/1538-4357/ac17ea](https://doi.org/10.3847/1538-4357/ac17ea). arXiv: [2105.11641](https://arxiv.org/abs/2105.11641) [astro-ph.HE].
- Abbott, R. et al. (June 2022a). “All-sky, all-frequency directional search for persistent gravitational waves from Advanced LIGO’s and Advanced Virgo’s first three observing runs”. In: *Physical Review D* 105.12, 122001, p. 122001. DOI: [10.1103/PhysRevD.105.122001](https://doi.org/10.1103/PhysRevD.105.122001). arXiv: [2110.09834](https://arxiv.org/abs/2110.09834) [gr-qc].
- Abbott, R. et al. (Nov. 2022b). “All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO and Advanced Virgo O3 data”. In: *Physical Review D* 106.10, 102008, p. 102008. DOI: [10.1103/PhysRevD.106.102008](https://doi.org/10.1103/PhysRevD.106.102008). arXiv: [2201.00697](https://arxiv.org/abs/2201.00697) [gr-qc].
- Abbott, R. et al. (Dec. 2022c). “Model-based Cross-correlation Search for Gravitational Waves from the Low-mass X-Ray Binary Scorpius X-1 in LIGO O3 Data”. In: *The Astrophysical Journal* 941.2, L30, p. L30. DOI: [10.3847/2041-8213/aca1b0](https://doi.org/10.3847/2041-8213/aca1b0). arXiv: [2209.02863](https://arxiv.org/abs/2209.02863) [astro-ph.HE].

- Abbott, R. et al. (June 2022d). “Narrowband Searches for Continuous and Long-duration Transient Gravitational Waves from Known Pulsars in the LIGO-Virgo Third Observing Run”. In: *The Astrophysical Journal* 932.2, 133, p. 133. DOI: [10.3847/1538-4357/ac6ad0](https://doi.org/10.3847/1538-4357/ac6ad0). arXiv: [2112.10990](https://arxiv.org/abs/2112.10990) [gr-qc].
- Abbott, R. et al. (Aug. 2022e). “Search for continuous gravitational wave emission from the Milky Way center in O3 LIGO-Virgo data”. In: *Physical Review D* 106.4, 042003, p. 042003. DOI: [10.1103/PhysRevD.106.042003](https://doi.org/10.1103/PhysRevD.106.042003). arXiv: [2204.04523](https://arxiv.org/abs/2204.04523) [astro-ph.HE].
- (Sept. 2022f). “Search for gravitational waves from Scorpius X-1 with a hidden Markov model in O3 LIGO data”. In: *Physical Review D* 106.6, 062002, p. 062002. DOI: [10.1103/PhysRevD.106.062002](https://doi.org/10.1103/PhysRevD.106.062002). arXiv: [2201.10104](https://arxiv.org/abs/2201.10104) [gr-qc].
- Abbott, R. et al. (Apr. 2022g). “Search of the early O3 LIGO data for continuous gravitational waves from the Cassiopeia A and Vela Jr. supernova remnants”. In: *Physical Review D* 105.8, 082005, p. 082005. DOI: [10.1103/PhysRevD.105.082005](https://doi.org/10.1103/PhysRevD.105.082005). arXiv: [2111.15116](https://arxiv.org/abs/2111.15116) [gr-qc].
- Abbott, R. et al. (Aug. 2022h). “Searches for Gravitational Waves from Known Pulsars at Two Harmonics in the Second and Third LIGO-Virgo Observing Runs”. In: *The Astrophysical Journal* 935.1, 1, p. 1. DOI: [10.3847/1538-4357/ac6acf](https://doi.org/10.3847/1538-4357/ac6acf). arXiv: [2111.13106](https://arxiv.org/abs/2111.13106) [astro-ph.HE].
- Abbott, R. et al. (Oct. 2023a). “GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo during the Second Part of the Third Observing Run”. In: *Physical Review X* 13.4, 041039, p. 041039. DOI: [10.1103/PhysRevX.13.041039](https://doi.org/10.1103/PhysRevX.13.041039). arXiv: [2111.03606](https://arxiv.org/abs/2111.03606) [gr-qc].
- Abbott, R. et al. (Aug. 2023b). “Open Data from the Third Observing Run of LIGO, Virgo, KAGRA, and GEO”. in: *The Astrophysical Journal Supplement Series* 267.2, 29, p. 29. DOI: [10.3847/1538-4365/acdc9f](https://doi.org/10.3847/1538-4365/acdc9f). arXiv: [2302.03676](https://arxiv.org/abs/2302.03676) [gr-qc].
- Abolmasov, Pavel, Anton Biryukov, and Sergei B. Popov (Feb. 2024). “Spin Evolution of Neutron Stars”. In: *Galaxies* 12.1, 7, p. 7. DOI: [10.3390/galaxies12010007](https://doi.org/10.3390/galaxies12010007). arXiv: [2402.04331](https://arxiv.org/abs/2402.04331) [astro-ph.HE].
- Acernese, F. et al. (Jan. 2015). “Advanced Virgo: a second-generation interferometric gravitational wave detector”. In: *Classical and Quantum Gravity* 32.2, 024001, p. 024001. DOI: [10.1088/0264-9381/32/2/024001](https://doi.org/10.1088/0264-9381/32/2/024001). arXiv: [1408.3978](https://arxiv.org/abs/1408.3978) [gr-qc].
- Agazie, Gabriella et al. (July 2023). “The NANOGrav 15 yr Data Set: Evidence for a Gravitational-wave Background”. In: *The Astrophysical Journal* 951.1, L8, p. L8. DOI: [10.3847/2041-8213/acdac6](https://doi.org/10.3847/2041-8213/acdac6). arXiv: [2306.16213](https://arxiv.org/abs/2306.16213) [astro-ph.HE].
- Akbal, O., M. A. Alpar, S. Buchner, and D. Pines (Aug. 2017). “Nonlinear interglitch dynamics, the braking index of the Vela pulsar and the time to the next glitch”. In: *Monthly Notices of the Royal Astronomical Society* 469.4, pp. 4183–4192. DOI: [10.1093/mnras/stx1095](https://doi.org/10.1093/mnras/stx1095). arXiv: [1612.03805](https://arxiv.org/abs/1612.03805) [astro-ph.HE].
- Alford, Mark G. and Gerald Good (July 2008). “Flux tubes and the type-I/type-II transition in a superconductor coupled to a superfluid”. In: *Physical Review B* 78.2, 024510, p. 024510. DOI: [10.1103/PhysRevB.78.024510](https://doi.org/10.1103/PhysRevB.78.024510). arXiv: [0712.1810](https://arxiv.org/abs/0712.1810) [nucl-th].
- Allan, David W. (1987). “Should the classical variance be used as a basic measure in standards metrology?” In: *IEEE Transactions on Instrumentation and Measurement* IM-36.2, pp. 646–654. DOI: [10.1109/TIM.1987.6312761](https://doi.org/10.1109/TIM.1987.6312761).

- Alpar, M. A., P. W. Anderson, D. Pines, and J. Shaham (Mar. 1984a). “Vortex creep and the internal temperature of neutron stars. II. VELA pulsar.” In: *The Astrophysical Journal* 278, pp. 791–805. DOI: [10.1086/161849](https://doi.org/10.1086/161849).
- Alpar, M. A., H. F. Chau, K. S. Cheng, and D. Pines (May 1994). “Postglitch Relaxation of the Crab Pulsar: Evidence for Crust Cracking”. In: *The Astrophysical Journal* 427, p. L29. DOI: [10.1086/187357](https://doi.org/10.1086/187357).
- (Mar. 1996). “Postglitch Relaxation of the Crab Pulsar after Its First Four Major Glitches: The Combined Effects of Crust Cracking, Formation of Vortex Depletion Region and Vortex Creep”. In: *The Astrophysical Journal* 459, p. 706. DOI: [10.1086/176935](https://doi.org/10.1086/176935).
- Alpar, M. A., A. F. Cheng, M. A. Ruderman, and J. Shaham (Dec. 1982). “A new class of radio pulsars”. In: *Nature* 300.5894, pp. 728–730. DOI: [10.1038/300728a0](https://doi.org/10.1038/300728a0).
- Alpar, M. A., K. S. Cheng, and D. Pines (Nov. 1989). “Vortex Creep and the Internal Temperature of Neutron Stars: Linear and Nonlinear Response to a Glitch”. In: *The Astrophysical Journal* 346, p. 823. DOI: [10.1086/168063](https://doi.org/10.1086/168063).
- Alpar, M. A., D. Pines, P. W. Anderson, and J. Shaham (Jan. 1984b). “Vortex creep and the internal temperature of neutron stars. I - General theory”. In: *The Astrophysical Journal* 276, pp. 325–334. DOI: [10.1086/161616](https://doi.org/10.1086/161616).
- Alpar, M. Ali, Radha Nandkumar, and David Pines (Dec. 1986). “Vortex creep and the internal temperature of neutron stars : timing noise in pulsars.” In: *The Astrophysical Journal* 311, pp. 197–213. DOI: [10.1086/164765](https://doi.org/10.1086/164765).
- Andersen, Bridget C. and Scott M. Ransom (Aug. 2018). “A Fourier Domain “Jerk” Search for Binary Pulsars”. In: *The Astrophysical Journal* 863.1, L13, p. L13. DOI: [10.3847/2041-8213/aad59f](https://doi.org/10.3847/2041-8213/aad59f). arXiv: [1807.07900](https://arxiv.org/abs/1807.07900) [astro-ph.HE].
- Anderson, P. W. and N. Itoh (July 1975). “Pulsar glitches and restlessness as a hard superfluidity phenomenon”. In: *Nature* 256.5512, pp. 25–27. DOI: [10.1038/256025a0](https://doi.org/10.1038/256025a0).
- Andersson, N., G. L. Comer, and R. Prix (Mar. 2003). “Are Pulsar Glitches Triggered by a Superfluid Two-Stream Instability?” In: *Physical Review Letters* 90.9, 091101, p. 091101. DOI: [10.1103/PhysRevLett.90.091101](https://doi.org/10.1103/PhysRevLett.90.091101). arXiv: [astro-ph/0210486](https://arxiv.org/abs/astro-ph/0210486) [astro-ph].
- Andersson, N., K. Glampedakis, W. C. G. Ho, and C. M. Espinoza (Dec. 2012). “Pulsar Glitches: The Crust is not Enough”. In: *Physical Review Letters* 109.24, 241103, p. 241103. DOI: [10.1103/PhysRevLett.109.241103](https://doi.org/10.1103/PhysRevLett.109.241103). arXiv: [1207.0633](https://arxiv.org/abs/1207.0633) [astro-ph.SR].
- Andersson, Nils (Aug. 1998). “A New Class of Unstable Modes of Rotating Relativistic Stars”. In: *The Astrophysical Journal* 502.2, pp. 708–713. DOI: [10.1086/305919](https://doi.org/10.1086/305919). arXiv: [gr-qc/9706075](https://arxiv.org/abs/gr-qc/9706075) [gr-qc].
- Andersson, Nils, Kostas D. Kokkotas, and Nikolaos Stergioulas (May 1999). “On the Relevance of the R-Mode Instability for Accreting Neutron Stars and White Dwarfs”. In: *The Astrophysical Journal* 516.1, pp. 307–314. DOI: [10.1086/307082](https://doi.org/10.1086/307082). arXiv: [astro-ph/9806089](https://arxiv.org/abs/astro-ph/9806089) [astro-ph].
- Antonelli, Marco, Avishek Basu, and Brynmor Haskell (Apr. 2023). “Stochastic processes for pulsar timing noise: fluctuations in the internal and external torques”. In: *Monthly Notices of the Royal Astronomical Society* 520.2, pp. 2813–2828. DOI: [10.1093/mnras/stad256](https://doi.org/10.1093/mnras/stad256). arXiv: [2206.10416](https://arxiv.org/abs/2206.10416) [astro-ph.HE].

- Antonopoulou, D., C. M. Espinoza, L. Kuiper, and N. Andersson (Jan. 2018). “Pulsar spin-down: the glitch-dominated rotation of PSR J0537-6910”. In: *Monthly Notices of the Royal Astronomical Society* 473.2, pp. 1644–1655. DOI: [10.1093/mnras/stx2429](https://doi.org/10.1093/mnras/stx2429). arXiv: [1708.09459](https://arxiv.org/abs/1708.09459) [astro-ph.HE].
- Antonov, Ivan and Mark Borodovsky (2010). “GeneTack: Frameshift identification in protein-coding sequences by the Viterbi algorithm”. In: *Journal of Bioinformatics and Computational Biology* 08.03, pp. 535–551. DOI: [10.1142/S0219720010004847](https://doi.org/10.1142/S0219720010004847).
- Archibald, Anne M. et al. (June 2009). “A Radio Pulsar/X-ray Binary Link”. In: *Science* 324.5933, p. 1411. DOI: [10.1126/science.1172740](https://doi.org/10.1126/science.1172740). arXiv: [0905.3397](https://arxiv.org/abs/0905.3397) [astro-ph.HE].
- Arras, Phil, Eanna E. Flanagan, Sharon M. Morsink, A. Katrin Schenk, Saul A. Teukolsky, and Ira Wasserman (July 2003). “Saturation of the r-Mode Instability”. In: *The Astrophysical Journal* 591.2, pp. 1129–1151. DOI: [10.1086/374657](https://doi.org/10.1086/374657). arXiv: [astro-ph/0202345](https://arxiv.org/abs/astro-ph/0202345) [astro-ph].
- Arvanitaki, Asimina, Savas Dimopoulos, Sergei Dubovsky, Nemanja Kaloper, and John March-Russell (June 2010). “String axiverse”. In: *Physical Review D* 81.12, 123530, p. 123530. DOI: [10.1103/PhysRevD.81.123530](https://doi.org/10.1103/PhysRevD.81.123530). arXiv: [0905.4720](https://arxiv.org/abs/0905.4720) [hep-th].
- Arzoumanian, Z., D. J. Nice, J. H. Taylor, and S. E. Thorsett (Feb. 1994). “Timing Behavior of 96 Radio Pulsars”. In: *The Astrophysical Journal* 422, p. 671. DOI: [10.1086/173760](https://doi.org/10.1086/173760).
- Ashok, A., B. Beheshtipour, M. A. Papa, P. C. C. Freire, B. Steltner, B. Machenschalk, O. Behnke, B. Allen, and R. Prix (Dec. 2021). “New Searches for Continuous Gravitational Waves from Seven Fast Pulsars”. In: *The Astrophysical Journal* 923.1, 85, p. 85. DOI: [10.3847/1538-4357/ac2582](https://doi.org/10.3847/1538-4357/ac2582). arXiv: [2107.09727](https://arxiv.org/abs/2107.09727) [astro-ph.HE].
- Ashton, G. and R. Prix (May 2018). “Hierarchical multistage MCMC follow-up of continuous gravitational wave candidates”. In: *Physical Review D* 97.10, 103020, p. 103020. DOI: [10.1103/PhysRevD.97.103020](https://doi.org/10.1103/PhysRevD.97.103020). arXiv: [1802.05450](https://arxiv.org/abs/1802.05450) [astro-ph.IM].
- Ashton, G., R. Prix, and D. I. Jones (Sept. 2017). “Statistical characterization of pulsar glitches and their potential impact on searches for continuous gravitational waves”. In: *Physical Review D* 96.6, 063004, p. 063004. DOI: [10.1103/PhysRevD.96.063004](https://doi.org/10.1103/PhysRevD.96.063004). arXiv: [1704.00742](https://arxiv.org/abs/1704.00742) [gr-qc].
- Ashton, Gregory, Paul D. Lasky, Vanessa Graber, and Jim Palfreyman (Aug. 2019a). “Rotational evolution of the Vela pulsar during the 2016 glitch”. In: *Nature Astronomy* 3, pp. 1143–1148. DOI: [10.1038/s41550-019-0844-6](https://doi.org/10.1038/s41550-019-0844-6). arXiv: [1907.01124](https://arxiv.org/abs/1907.01124) [astro-ph.HE].
- Ashton, Gregory et al. (Apr. 2019b). “BILBY: A User-friendly Bayesian Inference Library for Gravitational-wave Astronomy”. In: *The Astrophysical Journal Supplement Series* 241.2, 27, p. 27. DOI: [10.3847/1538-4365/ab06fc](https://doi.org/10.3847/1538-4365/ab06fc). arXiv: [1811.02042](https://arxiv.org/abs/1811.02042) [astro-ph.IM].
- Astone, P., A. Colla, S. D’Antonio, S. Frasca, and C. Palomba (June 2012). “Coherent search of continuous gravitational wave signals: extension of the 5-vectors method to a network of detectors”. In: *Journal of Physics Conference Series*. Vol. 363. Journal of Physics Conference Series. IOP, 012038, p. 012038. DOI: [10.1088/1742-6596/363/1/012038](https://doi.org/10.1088/1742-6596/363/1/012038). arXiv: [1203.6733](https://arxiv.org/abs/1203.6733) [astro-ph.IM].
- Astone, Pia, Alberto Colla, Sabrina D’Antonio, Sergio Frasca, and Cristiano Palomba (Aug. 2014). “Method for all-sky searches of continuous gravitational wave signals using the frequency-Hough transform”. In: *Physical Review D* 90.4, 042002, p. 042002. DOI: [10.1103/PhysRevD.90.042002](https://doi.org/10.1103/PhysRevD.90.042002). arXiv: [1407.8333](https://arxiv.org/abs/1407.8333) [astro-ph.IM].

- Atwood, W. B. et al. (June 2009). “The Large Area Telescope on the Fermi Gamma-Ray Space Telescope Mission”. In: *The Astrophysical Journal* 697.2, pp. 1071–1102. DOI: [10.1088/0004-637X/697/2/1071](https://doi.org/10.1088/0004-637X/697/2/1071). arXiv: [0902.1089](https://arxiv.org/abs/0902.1089) [astro-ph.IM].
- Backer, D. C. (Oct. 1970). “Pulsar Nulling Phenomena”. In: *Nature* 228.5266, pp. 42–43. DOI: [10.1038/228042a0](https://doi.org/10.1038/228042a0).
- Bailes, M. et al. (Sept. 2011). “Transformation of a Star into a Planet in a Millisecond Pulsar Binary”. In: *Science* 333.6050, p. 1717. DOI: [10.1126/science.1208890](https://doi.org/10.1126/science.1208890). arXiv: [1108.5201](https://arxiv.org/abs/1108.5201) [astro-ph.SR].
- Bailes, M. et al. (Oct. 2017). “The UTMOST: A Hybrid Digital Signal Processor Transforms the Molonglo Observatory Synthesis Telescope”. In: *Publications of the Astronomical Society of Australia* 34, e045, e045. DOI: [10.1017/pasa.2017.39](https://doi.org/10.1017/pasa.2017.39). arXiv: [1708.09619](https://arxiv.org/abs/1708.09619) [astro-ph.IM].
- Balakrishnan, Vishnu, David Champion, Ewan Barr, Michael Kramer, V. Venkatraman Krishnan, Ralph P. Eatough, Rahul Sengar, and Matthew Bailes (Mar. 2022). “Coherent search for binary pulsars across all Five Keplerian parameters in radio observations using the template-bank algorithm”. In: *Monthly Notices of the Royal Astronomical Society* 511.1, pp. 1265–1284. DOI: [10.1093/mnras/stab3746](https://doi.org/10.1093/mnras/stab3746). arXiv: [2112.11991](https://arxiv.org/abs/2112.11991) [astro-ph.IM].
- Banagiri, Sharan, Ling Sun, Michael W. Coughlin, and Andrew Melatos (July 2019). “Search strategies for long gravitational-wave transients: Hidden Markov model tracking and seedless clustering”. In: *Physical Review D* 100.2, 024034, p. 024034. DOI: [10.1103/PhysRevD.100.024034](https://doi.org/10.1103/PhysRevD.100.024034). arXiv: [1903.02638](https://arxiv.org/abs/1903.02638) [astro-ph.IM].
- Barboza Rossetto, Pedro Henrique, Jörg Frauendiener, and Andrew Melatos (Mar. 2024). “Quadrupole Moment of a Magnetically Confined Mountain on an Accreting Neutron Star in General Relativity”. In: *arXiv e-prints*, arXiv:2404.00339, arXiv:2404.00339. DOI: [10.48550/arXiv.2404.00339](https://doi.org/10.48550/arXiv.2404.00339). arXiv: [2404.00339](https://arxiv.org/abs/2404.00339) [astro-ph.HE].
- Basu, A., B. Shaw, D. Antonopoulou, M. J. Keith, A. G. Lyne, M. B. Mickaliger, B. W. Stappers, P. Weltevrede, and C. A. Jordan (Mar. 2022). “The Jodrell bank glitch catalogue: 106 new rotational glitches in 70 pulsars”. In: *Monthly Notices of the Royal Astronomical Society* 510.3, pp. 4049–4062. DOI: [10.1093/mnras/stab3336](https://doi.org/10.1093/mnras/stab3336). arXiv: [2111.06835](https://arxiv.org/abs/2111.06835) [astro-ph.HE].
- Basu, Avishek, Bhal Chandra Joshi, M. A. Krishnakumar, Dipankar Bhattacharya, Rana Nandi, Debades Bandhopadhyay, Prasanta Char, and P. K. Manoharan (Jan. 2020). “Observed glitches in eight young pulsars”. In: *Monthly Notices of the Royal Astronomical Society* 491.3, pp. 3182–3191. DOI: [10.1093/mnras/stz3230](https://doi.org/10.1093/mnras/stz3230). arXiv: [1911.04934](https://arxiv.org/abs/1911.04934) [astro-ph.HE].
- Baumgardt, H. (Jan. 2017). “N -body modelling of globular clusters: masses, mass-to-light ratios and intermediate-mass black holes”. In: *Monthly Notices of the Royal Astronomical Society* 464.2, pp. 2174–2202. DOI: [10.1093/mnras/stw2488](https://doi.org/10.1093/mnras/stw2488). arXiv: [1609.08794](https://arxiv.org/abs/1609.08794) [astro-ph.GA].
- Baykal, A. and H. Oegelman (Jan. 1993). “An empirical torque noise and spin-up model for accretion-powered X-ray pulsars.” In: *Astronomy and Astrophysics* 267, pp. 119–125.

- Bayley, Joe, Chris Messenger, and Graham Woan (July 2019). “Generalized application of the Viterbi algorithm to searches for continuous gravitational-wave signals”. In: *Physical Review D* 100.2, 023006, p. 023006. DOI: [10.1103/PhysRevD.100.023006](https://doi.org/10.1103/PhysRevD.100.023006). arXiv: [1903.12614](https://arxiv.org/abs/1903.12614) [astro-ph.IM].
- Baym, G. and D. Pines (Jan. 1971). “Neutron starquakes and pulsar speedup.” In: *Annals of Physics* 66, pp. 816–835. DOI: [10.1016/0003-4916\(71\)90084-4](https://doi.org/10.1016/0003-4916(71)90084-4).
- Baym, Gordon, Tetsuo Hatsuda, Toru Kojo, Philip D. Powell, Yifan Song, and Tatsuyuki Takatsuka (May 2018). “From hadrons to quarks in neutron stars: a review”. In: *Reports on Progress in Physics* 81.5, 056902, p. 056902. DOI: [10.1088/1361-6633/aaae14](https://doi.org/10.1088/1361-6633/aaae14). arXiv: [1707.04966](https://arxiv.org/abs/1707.04966) [astro-ph.HE].
- Baym, Gordon, Christopher Pethick, and David Pines (Nov. 1969a). “Superfluidity in Neutron Stars”. In: *Nature* 224.5220, pp. 673–674. DOI: [10.1038/224673a0](https://doi.org/10.1038/224673a0).
- Baym, Gordon, Christopher Pethick, David Pines, and Malvin Ruderman (Nov. 1969b). “Spin Up in Neutron Stars : The Future of the Vela Pulsar”. In: *Nature* 224.5222, pp. 872–874. DOI: [10.1038/224872a0](https://doi.org/10.1038/224872a0).
- Behrens, E. A. et al. (Apr. 2020). “The NANOGrav 11 yr Data Set: Constraints on Planetary Masses Around 45 Millisecond Pulsars”. In: *The Astrophysical Journal* 893.1, L8, p. L8. DOI: [10.3847/2041-8213/ab8121](https://doi.org/10.3847/2041-8213/ab8121). arXiv: [1912.00482](https://arxiv.org/abs/1912.00482) [astro-ph.EP].
- Beniwal, Deeksha, Patrick Clearwater, Liam Dunn, Andrew Melatos, and David Ottaway (Apr. 2021). “Search for continuous gravitational waves from ten H.E.S.S. sources using a hidden Markov model”. In: *Physical Review D* 103.8, 083009, p. 083009. DOI: [10.1103/PhysRevD.103.083009](https://doi.org/10.1103/PhysRevD.103.083009). arXiv: [2102.06334](https://arxiv.org/abs/2102.06334) [astro-ph.HE].
- Beniwal, Deeksha, Patrick Clearwater, Liam Dunn, Lucy Strang, Gavin Rowell, Andrew Melatos, and David Ottaway (Nov. 2022). “Search for continuous gravitational waves from HESS J1427-608 with a hidden Markov model”. In: *Physical Review D* 106.10, 103018, p. 103018. DOI: [10.1103/PhysRevD.106.103018](https://doi.org/10.1103/PhysRevD.106.103018). arXiv: [2210.09592](https://arxiv.org/abs/2210.09592) [astro-ph.HE].
- Bhattacharya, D. and G. Srinivasan (1991). “The Evolution of Neutron Star Magnetic Fields”. In: *Neutron Stars: Theory and Observation*. Ed. by Joseph Ventura and David Pines. Dordrecht: Springer Netherlands, pp. 219–233. ISBN: 978-94-011-3536-8. DOI: [10.1007/978-94-011-3536-8_14](https://doi.org/10.1007/978-94-011-3536-8_14).
- Biggs, J. D., M. Bailes, A. G. Lyne, W. M. Goss, and A. S. Fruchter (Mar. 1994). “Two radio pulsars in the globular cluster NGC 6624.” In: *Monthly Notices of the Royal Astronomical Society* 267, pp. 125–128. DOI: [10.1093/mnras/267.1.125](https://doi.org/10.1093/mnras/267.1.125).
- Bildsten, Lars (July 1998). “Gravitational Radiation and Rotation of Accreting Neutron Stars”. In: *The Astrophysical Journal* 501.1, pp. L89–L93. DOI: [10.1086/311440](https://doi.org/10.1086/311440). arXiv: [astro-ph/9804325](https://arxiv.org/abs/astro-ph/9804325) [astro-ph].
- Bildsten, Lars et al. (Dec. 1997). “Observations of Accreting Pulsars”. In: *The Astrophysical Journal Supplement Series* 113.2, pp. 367–408. DOI: [10.1086/313060](https://doi.org/10.1086/313060). arXiv: [astro-ph/9707125](https://arxiv.org/abs/astro-ph/9707125) [astro-ph].
- Blanchet, Luc (Dec. 2014). “Gravitational Radiation from Post-Newtonian Sources and Inspiralling Compact Binaries”. In: *Living Reviews in Relativity* 17.1, 2, p. 2. DOI: [10.12942/lrr-2014-2](https://doi.org/10.12942/lrr-2014-2). arXiv: [1310.1528](https://arxiv.org/abs/1310.1528) [gr-qc].
- Blandford, Roger D. and Roger W. Romani (Oct. 1988). “On the interpretation of pulsar braking indices.” In: *Monthly Notices of the Royal Astronomical Society* 234, 57P–60. DOI: [10.1093/mnras/234.1.57P](https://doi.org/10.1093/mnras/234.1.57P).

- Blaschke, David and Nicolas Chamel (Jan. 2018). “Phases of Dense Matter in Compact Stars”. In: *Astrophysics and Space Science Library*. Ed. by Luciano Rezzolla, Pierre Pizzochero, David Ian Jones, Nanda Rea, and Isaac Vidaña. Vol. 457. Astrophysics and Space Science Library, p. 337. DOI: [10.1007/978-3-319-97616-7_7](https://doi.org/10.1007/978-3-319-97616-7_7). arXiv: [1803.01836](https://arxiv.org/abs/1803.01836) [nucl-th].
- Bochenek, C. D., V. Ravi, K. V. Belov, G. Hallinan, J. Kocz, S. R. Kulkarni, and D. L. McKenna (Nov. 2020). “A fast radio burst associated with a Galactic magnetar”. In: *Nature* 587.7832, pp. 59–62. DOI: [10.1038/s41586-020-2872-x](https://doi.org/10.1038/s41586-020-2872-x). arXiv: [2005.10828](https://arxiv.org/abs/2005.10828) [astro-ph.HE].
- Bonazzola, S. and E. Gourgoulhon (Aug. 1996). “Gravitational waves from pulsars: emission by the magnetic-field-induced distortion.” In: *Astronomy and Astrophysics* 312, pp. 675–690. DOI: [10.48550/arXiv.astro-ph/9602107](https://doi.org/10.48550/arXiv.astro-ph/9602107). arXiv: [astro-ph/9602107](https://arxiv.org/abs/astro-ph/9602107) [astro-ph].
- Bondarescu, Ruxandra, Saul A. Teukolsky, and Ira Wasserman (Sept. 2007). “Spin evolution of accreting neutron stars: Nonlinear development of the r-mode instability”. In: *Physical Review D* 76.6, 064019, p. 064019. DOI: [10.1103/PhysRevD.76.064019](https://doi.org/10.1103/PhysRevD.76.064019). arXiv: [0704.0799](https://arxiv.org/abs/0704.0799) [astro-ph].
- Bondi, H. (May 1957). “Plane Gravitational Waves in General Relativity”. In: *Nature* 179.4569, pp. 1072–1073. DOI: [10.1038/1791072a0](https://doi.org/10.1038/1791072a0).
- Boyles, J., D. R. Lorimer, P. J. Turk, R. Mnatsakanov, R. S. Lynch, S. M. Ransom, P. C. Freire, and K. Belczynski (Nov. 2011). “Young Radio Pulsars in Galactic Globular Clusters”. In: *The Astrophysical Journal* 742.1, 51, p. 51. DOI: [10.1088/0004-637X/742/1/51](https://doi.org/10.1088/0004-637X/742/1/51). arXiv: [1108.4402](https://arxiv.org/abs/1108.4402) [astro-ph.SR].
- Boynton, P. E., E. J. Groth, D. P. Hutchinson, Jr. Nanos G. P., R. B. Partridge, and D. T. Wilkinson (July 1972). “Optical Timing of the Crab Pulsar, NP 0532”. In: *The Astrophysical Journal* 175, p. 217. DOI: [10.1086/151550](https://doi.org/10.1086/151550).
- Brady, Patrick, Giovanni Losurdo, and Hisaaki Shinkai (2022). “LIGO, VIRGO, and KAGRA as the international gravitational wave network”. In: *Handbook of gravitational wave astronomy*. Springer, pp. 1–21.
- Bransgrove, Ashley, Andrei M. Beloborodov, and Yuri Levin (July 2020). “A Quake Quenching the Vela Pulsar”. In: *The Astrophysical Journal* 897.2, 173, p. 173. DOI: [10.3847/1538-4357/ab93b7](https://doi.org/10.3847/1538-4357/ab93b7). arXiv: [2001.08658](https://arxiv.org/abs/2001.08658) [astro-ph.HE].
- Brito, Richard, Vitor Cardoso, and Paolo Pani (July 2015). “Black holes as particle detectors: evolution of superradiant instabilities”. In: *Classical and Quantum Gravity* 32.13, 134001, p. 134001. DOI: [10.1088/0264-9381/32/13/134001](https://doi.org/10.1088/0264-9381/32/13/134001). arXiv: [1411.0686](https://arxiv.org/abs/1411.0686) [gr-qc].
- Brook, P. R., A. Karastergiou, S. Buchner, S. J. Roberts, M. J. Keith, S. Johnston, and R. M. Shannon (Jan. 2014). “Evidence of an Asteroid Encountering a Pulsar”. In: *The Astrophysical Journal* 780.2, L31, p. L31. DOI: [10.1088/2041-8205/780/2/L31](https://doi.org/10.1088/2041-8205/780/2/L31). arXiv: [1311.3541](https://arxiv.org/abs/1311.3541) [astro-ph.HE].
- Brook, P. R., A. Karastergiou, S. Johnston, M. Kerr, R. M. Shannon, and S. J. Roberts (Feb. 2016). “Emission-rotation correlation in pulsars: new discoveries with optimal techniques”. In: *Monthly Notices of the Royal Astronomical Society* 456.2, pp. 1374–1393. DOI: [10.1093/mnras/stv2715](https://doi.org/10.1093/mnras/stv2715). arXiv: [1511.05481](https://arxiv.org/abs/1511.05481) [astro-ph.HE].

- Bucciantini, N., T. A. Thompson, J. Arons, E. Quataert, and L. Del Zanna (June 2006). “Relativistic magnetohydrodynamics winds from rotating neutron stars”. In: *Monthly Notices of the Royal Astronomical Society* 368.4, pp. 1717–1734. DOI: [10.1111/j.1365-2966.2006.10217.x](https://doi.org/10.1111/j.1365-2966.2006.10217.x). arXiv: [astro-ph/0602475](https://arxiv.org/abs/astro-ph/0602475) [astro-ph].
- Buikema, A. et al. (Sept. 2020). “Sensitivity and performance of the Advanced LIGO detectors in the third observing run”. In: *Physical Review D* 102.6, 062003, p. 062003. DOI: [10.1103/PhysRevD.102.062003](https://doi.org/10.1103/PhysRevD.102.062003). arXiv: [2008.01301](https://arxiv.org/abs/2008.01301) [astro-ph.IM].
- Bult, Peter, Tod E. Strohmayer, Christian Malacaria, Mason Ng, and Zorawar Wadiasingh (May 2021). “Long-term Coherent Timing of the Accreting Millisecond Pulsar IGR J17062-6143”. In: *The Astrophysical Journal* 912.2, 120, p. 120. DOI: [10.3847/1538-4357/abf13f](https://doi.org/10.3847/1538-4357/abf13f). arXiv: [2103.12556](https://arxiv.org/abs/2103.12556) [astro-ph.HE].
- Caballero, R. N. et al. (Apr. 2016). “The noise properties of 42 millisecond pulsars from the European Pulsar Timing Array and their impact on gravitational-wave searches”. In: *Monthly Notices of the Royal Astronomical Society* 457.4, pp. 4421–4440. DOI: [10.1093/mnras/stw179](https://doi.org/10.1093/mnras/stw179). arXiv: [1510.09194](https://arxiv.org/abs/1510.09194) [astro-ph.IM].
- Camilo, F. and F. A. Rasio (July 2005). “Pulsars in Globular Clusters”. In: *Binary Radio Pulsars*. Ed. by Fred A. Rasio and Ingrid H. Stairs. Vol. 328. Astronomical Society of the Pacific Conference Series, p. 147. DOI: [10.48550/arXiv.astro-ph/0501226](https://doi.org/10.48550/arXiv.astro-ph/0501226). arXiv: [astro-ph/0501226](https://arxiv.org/abs/astro-ph/0501226) [astro-ph].
- Camilo, F. et al. (Apr. 2018). “Revival of the Magnetar PSR J1622-4950: Observations with MeerKAT, Parkes, XMM-Newton, Swift, Chandra, and NuSTAR”. in: *The Astrophysical Journal* 856.2, 180, p. 180. DOI: [10.3847/1538-4357/aab35a](https://doi.org/10.3847/1538-4357/aab35a). arXiv: [1804.01933](https://arxiv.org/abs/1804.01933) [astro-ph.HE].
- Camilo, Fernando, Scott M. Ransom, Jules P. Halpern, John Reynolds, David J. Helfand, Neil Zimmerman, and John Sarkissian (Aug. 2006). “Transient pulsed radio emission from a magnetar”. In: *Nature* 442.7105, pp. 892–895. DOI: [10.1038/nature04986](https://doi.org/10.1038/nature04986). arXiv: [astro-ph/0605429](https://arxiv.org/abs/astro-ph/0605429) [astro-ph].
- Caplan, M. E. and C. J. Horowitz (Oct. 2017). “Colloquium: Astromaterial science and nuclear pasta”. In: *Reviews of Modern Physics* 89.4, 041002, p. 041002. DOI: [10.1103/RevModPhys.89.041002](https://doi.org/10.1103/RevModPhys.89.041002). arXiv: [1606.03646](https://arxiv.org/abs/1606.03646) [astro-ph.HE].
- Caride, Santiago, Ra Inta, Benjamin J. Owen, and Binod Rajbhandari (Sept. 2019). “How to search for gravitational waves from r -modes of known pulsars”. In: *Physical Review D* 100.6, 064013, p. 064013. DOI: [10.1103/PhysRevD.100.064013](https://doi.org/10.1103/PhysRevD.100.064013). arXiv: [1907.04946](https://arxiv.org/abs/1907.04946) [gr-qc].
- Carlin, J. B. and A. Melatos (Oct. 2019a). “Autocorrelations in pulsar glitch waiting times and sizes”. In: *Monthly Notices of the Royal Astronomical Society* 488.4, pp. 4890–4896. DOI: [10.1093/mnras/stz2014](https://doi.org/10.1093/mnras/stz2014). arXiv: [1907.09143](https://arxiv.org/abs/1907.09143) [astro-ph.HE].
- (Mar. 2019b). “Generating quasi-periodic pulsar glitches using a state-dependent Poisson process”. In: *Monthly Notices of the Royal Astronomical Society* 483.4, pp. 4742–4750. DOI: [10.1093/mnras/sty3433](https://doi.org/10.1093/mnras/sty3433). arXiv: [1901.01633](https://arxiv.org/abs/1901.01633) [astro-ph.HE].
- (May 2020). “Long-term statistics of pulsar glitches triggered by a Brownian stress accumulation process”. In: *Monthly Notices of the Royal Astronomical Society* 494.3, pp. 3383–3391. DOI: [10.1093/mnras/staa935](https://doi.org/10.1093/mnras/staa935). arXiv: [2004.00168](https://arxiv.org/abs/2004.00168) [astro-ph.HE].
- Carlin, J. B., A. Melatos, and D. Vukcevic (Jan. 2019). “Temporal clustering of rotational glitches in the Crab pulsar”. In: *Monthly Notices of the Royal Astronomical Society* 482.3, pp. 3736–3743. DOI: [10.1093/mnras/sty2865](https://doi.org/10.1093/mnras/sty2865). arXiv: [1810.12458](https://arxiv.org/abs/1810.12458) [astro-ph.HE].

- Carlin, Julian B. and Andrew Melatos (Aug. 2021). “Long-term Statistics of Pulsar Glitches Due to History-dependent Avalanches”. In: *The Astrophysical Journal* 917.1, 1, p. 1. DOI: [10.3847/1538-4357/ac06a2](https://doi.org/10.3847/1538-4357/ac06a2). arXiv: [2105.13588](https://arxiv.org/abs/2105.13588) [astro-ph.HE].
- Carlin, Julian B. and Andrew Melatos (submitted). “How much spin wandering can continuous gravitational wave search algorithms handle?”
- Chamel, N. (Sept. 2017). “Superfluidity and Superconductivity in Neutron Stars”. In: *Journal of Astrophysics and Astronomy* 38.3, 43, p. 43. DOI: [10.1007/s12036-017-9470-9](https://doi.org/10.1007/s12036-017-9470-9). arXiv: [1709.07288](https://arxiv.org/abs/1709.07288) [astro-ph.HE].
- Chamel, Nicolas and Pawel Haensel (Dec. 2008). “Physics of Neutron Star Crusts”. In: *Living Reviews in Relativity* 11.1, 10, p. 10. DOI: [10.12942/lrr-2008-10](https://doi.org/10.12942/lrr-2008-10). arXiv: [0812.3955](https://arxiv.org/abs/0812.3955) [astro-ph].
- Chandler, Adam M. (Aug. 2003). “Pulsar searches: From radio to gamma-rays”. PhD thesis. California Institute of Technology.
- Chen, W. et al. (Apr. 2023). “MeerKAT discovery of 13 new pulsars in Omega Centauri”. In: *Monthly Notices of the Royal Astronomical Society* 520.3, pp. 3847–3856. DOI: [10.1093/mnras/stad029](https://doi.org/10.1093/mnras/stad029). arXiv: [2301.03864](https://arxiv.org/abs/2301.03864) [astro-ph.HE].
- Cheng, K. S. (Oct. 1987a). “Could Glitches Inducing Magnetospheric Fluctuations Produce Low-Frequency Pulsar Timing Noise?” In: *The Astrophysical Journal* 321, p. 805. DOI: [10.1086/165673](https://doi.org/10.1086/165673).
- (Oct. 1987b). “Outer Magnetospheric Fluctuations and Pulsar Timing Noise”. In: *The Astrophysical Journal* 321, p. 799. DOI: [10.1086/165672](https://doi.org/10.1086/165672).
- CHIME/FRB Collaboration et al. (Nov. 2020). “A bright millisecond-duration radio burst from a Galactic magnetar”. In: *Nature* 587.7832, pp. 54–58. DOI: [10.1038/s41586-020-2863-y](https://doi.org/10.1038/s41586-020-2863-y). arXiv: [2005.10324](https://arxiv.org/abs/2005.10324) [astro-ph.HE].
- CHIME/Pulsar Collaboration et al. (July 2021). “The CHIME Pulsar Project: System Overview”. In: *The Astrophysical Journal Supplement Series* 255.1, 5, p. 5. DOI: [10.3847/1538-4365/abfdcb](https://doi.org/10.3847/1538-4365/abfdcb). arXiv: [2008.05681](https://arxiv.org/abs/2008.05681) [astro-ph.IM].
- Chugunov, A. I. and C. J. Horowitz (Sept. 2010). “Breaking stress of neutron star crust”. In: *Monthly Notices of the Royal Astronomical Society* 407.1, pp. L54–L58. DOI: [10.1111/j.1745-3933.2010.00903.x](https://doi.org/10.1111/j.1745-3933.2010.00903.x). arXiv: [1006.2279](https://arxiv.org/abs/1006.2279) [astro-ph.SR].
- Chukwude, A. E. and Finbarr Chidi Odo (Oct. 2016). “Assessing the effects of timing irregularities on radio pulsars anomalous braking indices”. In: *Research in Astronomy and Astrophysics* 16.10, 150, p. 150. DOI: [10.1088/1674-4527/16/10/150](https://doi.org/10.1088/1674-4527/16/10/150).
- Chukwude, A. E. and J. O. Urama (Aug. 2010). “Observations of microglitches in Hartebeesthoek Radio Astronomy Observatory radio pulsars”. In: *Monthly Notices of the Royal Astronomical Society* 406.3, pp. 1907–1917. DOI: [10.1111/j.1365-2966.2010.16789.x](https://doi.org/10.1111/j.1365-2966.2010.16789.x). arXiv: [1005.1534](https://arxiv.org/abs/1005.1534) [astro-ph.SR].
- Clark, G. W. (Aug. 1975). “X-ray binaries in globular clusters.” In: *The Astrophysical Journal* 199, pp. L143–L145. DOI: [10.1086/181869](https://doi.org/10.1086/181869).
- Cognard, Ismaël and Donald C. Backer (Sept. 2004). “A Microglitch in the Millisecond Pulsar PSR B1821-24 in M28”. In: *The Astrophysical Journal* 612.2, pp. L125–L127. DOI: [10.1086/424692](https://doi.org/10.1086/424692). arXiv: [astro-ph/0407546](https://arxiv.org/abs/astro-ph/0407546) [astro-ph].
- Coles, W., G. Hobbs, D. J. Champion, R. N. Manchester, and J. P. W. Verbiest (Nov. 2011). “Pulsar timing analysis in the presence of correlated noise”. In: *Monthly Notices of the Royal Astronomical Society* 418.1, pp. 561–570. DOI: [10.1111/j.1365-2966.2011.19505.x](https://doi.org/10.1111/j.1365-2966.2011.19505.x). arXiv: [1107.5366](https://arxiv.org/abs/1107.5366) [astro-ph.IM].

- Contopoulos, I. and A. Spitkovsky (June 2006). “Revised Pulsar Spin-down”. In: *The Astrophysical Journal* 643.2, pp. 1139–1145. DOI: [10.1086/501161](https://doi.org/10.1086/501161). arXiv: [astro-ph/0512002](https://arxiv.org/abs/astro-ph/0512002) [astro-ph].
- Cordes, J. M. and G. S. Downs (Nov. 1985). “JPL pulsar timing observations. III. Pulsar rotation fluctuations.” In: *The Astrophysical Journal Supplement Series* 59, pp. 343–382. DOI: [10.1086/191076](https://doi.org/10.1086/191076).
- Cordes, J. M., G. S. Downs, and J. Krause-Polstorff (July 1988). “JPL Pulsar Timing Observations. V. Macro- and Microjumps in the VELA Pulsar 0833-45”. In: *The Astrophysical Journal* 330, p. 847. DOI: [10.1086/166518](https://doi.org/10.1086/166518).
- Cordes, J. M. and D. J. Helfand (July 1980). “Pulsar timing .III. Timing noise of 50 pulsars.” In: *The Astrophysical Journal* 239, pp. 640–650. DOI: [10.1086/158150](https://doi.org/10.1086/158150).
- Cordes, J. M. and R. M. Shannon (Aug. 2008). “Rocking the Lighthouse: Circumpulsar Asteroids and Radio Intermittency”. In: *The Astrophysical Journal* 682.2, pp. 1152–1165. DOI: [10.1086/589425](https://doi.org/10.1086/589425). arXiv: [astro-ph/0605145](https://arxiv.org/abs/astro-ph/0605145) [astro-ph].
- Covas, P. B. and Alicia M. Sintes (June 2019). “New method to search for continuous gravitational waves from unknown neutron stars in binary systems”. In: *Physical Review D* 99.12, 124019, p. 124019. DOI: [10.1103/PhysRevD.99.124019](https://doi.org/10.1103/PhysRevD.99.124019). arXiv: [1904.04873](https://arxiv.org/abs/1904.04873) [astro-ph.IM].
- Covas, P. B. et al. (Apr. 2018). “Identification and mitigation of narrow spectral artifacts that degrade searches for persistent gravitational waves in the first two observing runs of Advanced LIGO”. in: *Physical Review D* 97.8, 082002, p. 082002. DOI: [10.1103/PhysRevD.97.082002](https://doi.org/10.1103/PhysRevD.97.082002). arXiv: [1801.07204](https://arxiv.org/abs/1801.07204) [astro-ph.IM].
- Crowther, Paul (Aug. 2012). “Birth, life and death of massive stars”. In: *Astronomy and Geophysics* 53.4, pp. 4.30–4.36. DOI: [10.1111/j.1468-4004.2012.53430.x](https://doi.org/10.1111/j.1468-4004.2012.53430.x).
- Cutler, Curt (Oct. 2002). “Gravitational waves from neutron stars with large toroidal B fields”. In: *Physical Review D* 66.8, 084025, p. 084025. DOI: [10.1103/PhysRevD.66.084025](https://doi.org/10.1103/PhysRevD.66.084025). arXiv: [gr-qc/0206051](https://arxiv.org/abs/gr-qc/0206051) [gr-qc].
- Cutler, Curt and Bernard F. Schutz (Sept. 2005). “Generalized F-statistic: Multiple detectors and multiple gravitational wave pulsars”. In: *Physical Review D* 72.6, 063006, p. 063006. DOI: [10.1103/PhysRevD.72.063006](https://doi.org/10.1103/PhysRevD.72.063006). arXiv: [gr-qc/0504011](https://arxiv.org/abs/gr-qc/0504011) [gr-qc].
- D’Alessandro, F., P. M. McCulloch, P. A. Hamilton, and A. A. Deshpande (Dec. 1995). “The timing noise of 45 southern pulsars”. In: *Monthly Notices of the Royal Astronomical Society* 277.3, pp. 1033–1046. DOI: [10.1093/mnras/277.3.1033](https://doi.org/10.1093/mnras/277.3.1033).
- Dang, S. J., J. P. Yuan, R. N. Manchester, L. Li, N. Wang, J. B. Wang, G. Hobbs, Z. Y. Liu, and F. F. Kou (June 2020). “Results of 12 yr of Pulsar Timing at Nanshan. I.”. In: *The Astrophysical Journal* 896.2, 140, p. 140. DOI: [10.3847/1538-4357/ab9082](https://doi.org/10.3847/1538-4357/ab9082). arXiv: [2005.02200](https://arxiv.org/abs/2005.02200) [astro-ph.HE].
- Davis, D. et al. (July 2021). “LIGO detector characterization in the second and third observing runs”. In: *Classical and Quantum Gravity* 38.13, 135014, p. 135014. DOI: [10.1088/1361-6382/abfd85](https://doi.org/10.1088/1361-6382/abfd85). arXiv: [2101.11673](https://arxiv.org/abs/2101.11673) [astro-ph.IM].
- Davis, Derek, Thomas Massinger, Andrew Lundgren, Jennifer C. Driggers, Alex L. Urban, and Laura Nuttall (Mar. 2019). “Improving the sensitivity of Advanced LIGO using noise subtraction”. In: *Classical and Quantum Gravity* 36.5, 055011, p. 055011. DOI: [10.1088/1361-6382/ab01c5](https://doi.org/10.1088/1361-6382/ab01c5). arXiv: [1809.05348](https://arxiv.org/abs/1809.05348) [astro-ph.IM].
- Day, Cherie K. (2022). “Pinpointing the Origins of Fast Radios Bursts”. PhD thesis. Swinburne University of Technology.

- De, Kishalay and Yashwant Gupta (Feb. 2016). “A real-time coherent dedispersion pipeline for the giant metrewave radio telescope”. In: *Experimental Astronomy* 41.1-2, pp. 67–93. DOI: [10.1007/s10686-015-9476-8](https://doi.org/10.1007/s10686-015-9476-8). arXiv: [1509.00186](https://arxiv.org/abs/1509.00186) [astro-ph.IM].
- Deller, Adam and Chris Flynn (Mar. 2020). “Vintage telescope rebooted in the hunt for FRBs”. In: *Nature Astronomy* 4, pp. 292–292. DOI: [10.1038/s41550-020-1046-y](https://doi.org/10.1038/s41550-020-1046-y).
- Dergachev, V. (Oct. 2010). “On blind searches for noise dominated signals: a loosely coherent approach”. In: *Classical and Quantum Gravity* 27.20, 205017, p. 205017. DOI: [10.1088/0264-9381/27/20/205017](https://doi.org/10.1088/0264-9381/27/20/205017). arXiv: [1003.2178](https://arxiv.org/abs/1003.2178) [gr-qc].
- (Mar. 2012). “Loosely coherent searches for sets of well-modeled signals”. In: *Physical Review D* 85.6, 062003, p. 062003. DOI: [10.1103/PhysRevD.85.062003](https://doi.org/10.1103/PhysRevD.85.062003). arXiv: [1110.3297](https://arxiv.org/abs/1110.3297) [gr-qc].
- Dergachev, Vladimir and Maria Alessandra Papa (Oct. 2020). “Results from the First All-Sky Search for Continuous Gravitational Waves from Small-Ellipticity Sources”. In: *Physical Review Letters* 125.17, 171101, p. 171101. DOI: [10.1103/PhysRevLett.125.171101](https://doi.org/10.1103/PhysRevLett.125.171101). arXiv: [2004.08334](https://arxiv.org/abs/2004.08334) [gr-qc].
- (Aug. 2021). “Search for continuous gravitational waves from small-ellipticity sources at low frequencies”. In: *Physical Review D* 104.4, 043003, p. 043003. DOI: [10.1103/PhysRevD.104.043003](https://doi.org/10.1103/PhysRevD.104.043003). arXiv: [2104.09007](https://arxiv.org/abs/2104.09007) [gr-qc].
- (Apr. 2023). “Frequency-Resolved Atlas of the Sky in Continuous Gravitational Waves”. In: *Physical Review X* 13.2, 021020, p. 021020. DOI: [10.1103/PhysRevX.13.021020](https://doi.org/10.1103/PhysRevX.13.021020). arXiv: [2202.10598](https://arxiv.org/abs/2202.10598) [gr-qc].
- Dergachev, Vladimir, Maria Alessandra Papa, Benjamin Steltner, and Heinz-Bernd Eggenstein (Apr. 2019). “Loosely coherent search in LIGO O1 data for continuous gravitational waves from Terzan 5 and the Galactic Center”. In: *Physical Review D* 99.8, 084048, p. 084048. DOI: [10.1103/PhysRevD.99.084048](https://doi.org/10.1103/PhysRevD.99.084048). arXiv: [1903.02389](https://arxiv.org/abs/1903.02389) [gr-qc].
- Detweiler, S. (Dec. 1979). “Pulsar timing measurements and the search for gravitational waves”. In: *The Astrophysical Journal* 234, pp. 1100–1104. DOI: [10.1086/157593](https://doi.org/10.1086/157593).
- Dewey, R. J. and J. M. Cordes (Jan. 1989). “The scaling of radio pulsar timing noise with spin parameters”. In: *Timing Neutron Stars*. Ed. by H. Ögelman and E. P. J. van den Heuvel. Vol. 262. NATO Advanced Study Institute (ASI) Series C, p. 119. DOI: [10.1007/978-94-009-2273-0_7](https://doi.org/10.1007/978-94-009-2273-0_7).
- Dodson, R. G., P. M. McCulloch, and D. R. Lewis (Jan. 2002). “High Time Resolution Observations of the January 2000 Glitch in the Vela Pulsar”. In: *The Astrophysical Journal* 564.2, pp. L85–L88. DOI: [10.1086/339068](https://doi.org/10.1086/339068). arXiv: [astro-ph/0201005](https://arxiv.org/abs/astro-ph/0201005) [astro-ph].
- Dodson, Richard, Dion Lewis, and Peter McCulloch (Apr. 2007). “Two decades of pulsar timing of Vela”. In: *Astrophysics and Space Science* 308.1-4, pp. 585–589. DOI: [10.1007/s10509-007-9372-4](https://doi.org/10.1007/s10509-007-9372-4). arXiv: [astro-ph/0612371](https://arxiv.org/abs/astro-ph/0612371) [astro-ph].
- Drake, F. D. and Jr. Craft H. D. (May 1968). “Pulse Structure of Four Pulsars”. In: *Science* 160.3829, pp. 758–760. DOI: [10.1126/science.160.3829.758](https://doi.org/10.1126/science.160.3829.758).
- Driggers, J. C. et al. (Feb. 2019). “Improving astrophysical parameter estimation via offline noise subtraction for Advanced LIGO”. In: *Physical Review D* 99.4, 042001, p. 042001. DOI: [10.1103/PhysRevD.99.042001](https://doi.org/10.1103/PhysRevD.99.042001). arXiv: [1806.00532](https://arxiv.org/abs/1806.00532) [astro-ph.IM].
- Drummond, L. V. and A. Melatos (Dec. 2017). “Stability of interlinked neutron vortex and proton flux tube arrays in a neutron star: equilibrium configurations”. In: *Monthly Notices of the Royal Astronomical Society* 472.4, pp. 4851–4869. DOI: [10.1093/mnras/stx2301](https://doi.org/10.1093/mnras/stx2301). arXiv: [1709.02254](https://arxiv.org/abs/1709.02254) [astro-ph.HE].

- Drummond, L. V. and A. Melatos (Mar. 2018). “Stability of interlinked neutron vortex and proton flux-tube arrays in a neutron star - II. Far-from-equilibrium dynamics”. In: *Monthly Notices of the Royal Astronomical Society* 475.1, pp. 910–920. DOI: [10.1093/mnras/stx3197](https://doi.org/10.1093/mnras/stx3197). arXiv: [1712.02938](https://arxiv.org/abs/1712.02938) [astro-ph.HE].
- Dunn, L., M. E. Lower, and A. Melatos (July 2021a). “Effects of periodicity in observation scheduling on parameter estimation of pulsar glitches”. In: *Monthly Notices of the Royal Astronomical Society* 504.3, pp. 3399–3411. DOI: [10.1093/mnras/stab1097](https://doi.org/10.1093/mnras/stab1097). arXiv: [2104.07363](https://arxiv.org/abs/2104.07363) [astro-ph.HE].
- Dunn, L., A. Melatos, C. M. Espinoza, D. Antonopoulou, and R. Dodson (July 2023a). “A new small glitch in Vela discovered with a hidden Markov model”. In: *Monthly Notices of the Royal Astronomical Society* 522.4, pp. 5469–5478. DOI: [10.1093/mnras/stad1335](https://doi.org/10.1093/mnras/stad1335). arXiv: [2304.13382](https://arxiv.org/abs/2304.13382) [astro-ph.HE].
- Dunn, L. et al. (July 2021b). “Confirmation of glitch event observed in the Vela pulsar (PSR J0835-4510)”. In: *The Astronomer’s Telegram* 14807, p. 1.
- Dunn, L. et al. (Sept. 2022a). “Confirmation of glitch event observed in PSR J0742-2822”. In: *The Astronomer’s Telegram* 15631, p. 1.
- Dunn, L. et al. (May 2022b). “Systematic upper limits on the size of missing pulsar glitches in the first UTMOST open data release”. In: *Monthly Notices of the Royal Astronomical Society* 512.1, pp. 1469–1482. DOI: [10.1093/mnras/stac551](https://doi.org/10.1093/mnras/stac551). arXiv: [2202.12442](https://arxiv.org/abs/2202.12442) [astro-ph.HE].
- Dunn, L. et al. (Jan. 2023b). “Confirmation of glitch event observed in PSR J1740-3015”. In: *The Astronomer’s Telegram* 15839, p. 1.
- Dunn, Liam, Patrick Clearwater, Andrew Melatos, and Karl Wette (Feb. 2022c). “Graphics processing unit implementation of the $\{F\}$ -statistic for continuous gravitational wave searches”. In: *Classical and Quantum Gravity* 39.4, 045003, p. 045003. DOI: [10.1088/1361-6382/ac4616](https://doi.org/10.1088/1361-6382/ac4616). arXiv: [2201.00451](https://arxiv.org/abs/2201.00451) [gr-qc].
- Edwards, R. T., G. B. Hobbs, and R. N. Manchester (Nov. 2006). “TEMPO2, a new pulsar timing package - II. The timing model and precision estimates”. In: *Monthly Notices of the Royal Astronomical Society* 372.4, pp. 1549–1574. DOI: [10.1111/j.1365-2966.2006.10870.x](https://doi.org/10.1111/j.1365-2966.2006.10870.x). arXiv: [astro-ph/0607664](https://arxiv.org/abs/astro-ph/0607664) [astro-ph].
- Einstein, Albert (Jan. 1918). “Über Gravitationswellen”. In: *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften*, pp. 154–167.
- Ellis, Justin A., Michele Vallisneri, Stephen R. Taylor, and Paul T. Baker (Sept. 2020). *ENTERPRISE: Enhanced Numerical Toolbox Enabling a Robust Pulsar Inference Suite*. version v3.0.0. DOI: [10.5281/zenodo.4059815](https://doi.org/10.5281/zenodo.4059815).
- EPTA Collaboration et al. (Oct. 2023). “The second data release from the European Pulsar Timing Array. III. Search for gravitational wave signals”. In: *Astronomy and Astrophysics* 678, A50, A50. DOI: [10.1051/0004-6361/202346844](https://doi.org/10.1051/0004-6361/202346844). arXiv: [2306.16214](https://arxiv.org/abs/2306.16214) [astro-ph.HE].
- Espinoza, C. M., D. Antonopoulou, R. Dodson, M. Stepanova, and A. Scherer (Mar. 2021). “Small glitches and other rotational irregularities of the Vela pulsar”. In: *Astronomy and Astrophysics* 647, A25, A25. DOI: [10.1051/0004-6361/202039044](https://doi.org/10.1051/0004-6361/202039044). arXiv: [2007.02921](https://arxiv.org/abs/2007.02921) [astro-ph.HE].

- Espinoza, C. M., D. Antonopoulou, B. W. Stappers, A. Watts, and A. G. Lyne (May 2014). “Neutron star glitches have a substantial minimum size”. In: *Monthly Notices of the Royal Astronomical Society* 440.3, pp. 2755–2762. DOI: [10.1093/mnras/stu395](https://doi.org/10.1093/mnras/stu395). arXiv: [1402.7219](https://arxiv.org/abs/1402.7219) [astro-ph.HE].
- Espinoza, C. M., A. G. Lyne, and B. W. Stappers (Apr. 2017). “New long-term braking index measurements for glitching pulsars using a glitch-template method”. In: *Monthly Notices of the Royal Astronomical Society* 466.1, pp. 147–162. DOI: [10.1093/mnras/stw3081](https://doi.org/10.1093/mnras/stw3081). arXiv: [1611.08314](https://arxiv.org/abs/1611.08314) [astro-ph.HE].
- Espinoza, C. M., A. G. Lyne, B. W. Stappers, and M. Kramer (June 2011). “A study of 315 glitches in the rotation of 102 pulsars”. In: *Monthly Notices of the Royal Astronomical Society* 414.2, pp. 1679–1704. DOI: [10.1111/j.1365-2966.2011.18503.x](https://doi.org/10.1111/j.1365-2966.2011.18503.x). arXiv: [1102.1743](https://arxiv.org/abs/1102.1743) [astro-ph.HE].
- Evans, Matthew et al. (June 2023). “Cosmic Explorer: A Submission to the NSF MPSAC ngGW Subcommittee”. In: *arXiv e-prints*, arXiv:2306.13745, arXiv:2306.13745. DOI: [10.48550/arXiv.2306.13745](https://doi.org/10.48550/arXiv.2306.13745). arXiv: [2306.13745](https://arxiv.org/abs/2306.13745) [astro-ph.IM].
- Eya, Innocent Okwudili, Johnson Ozoemene Urama, and Augustine Ejikeme Chukwude (June 2019). “On the distributions of pulsar glitch sizes and the inter-glitch time intervals”. In: *Research in Astronomy and Astrophysics* 19.6, 089, p. 089. DOI: [10.1088/1674-4527/19/6/89](https://doi.org/10.1088/1674-4527/19/6/89). arXiv: [1812.10053](https://arxiv.org/abs/1812.10053) [astro-ph.HE].
- Fattoyev, F. J., C. J. Horowitz, and Hao Lu (Apr. 2018). “Crust breaking and the limiting rotational frequency of neutron stars”. In: *arXiv e-prints*, arXiv:1804.04952, arXiv:1804.04952. DOI: [10.48550/arXiv.1804.04952](https://doi.org/10.48550/arXiv.1804.04952). arXiv: [1804.04952](https://arxiv.org/abs/1804.04952) [astro-ph.HE].
- Ferdman, R. D., R. F. Archibald, K. N. Gourgouliatos, and V. M. Kaspi (Jan. 2018). “The Glitches and Rotational History of the Highly Energetic Young Pulsar PSR J0537-6910”. In: *The Astrophysical Journal* 852.2, 123, p. 123. DOI: [10.3847/1538-4357/aaa198](https://doi.org/10.3847/1538-4357/aaa198). arXiv: [1708.08876](https://arxiv.org/abs/1708.08876) [astro-ph.HE].
- Feroz, F., M. P. Hobson, and M. Bridges (Oct. 2009). “MULTINEST: an efficient and robust Bayesian inference tool for cosmology and particle physics”. In: *Monthly Notices of the Royal Astronomical Society* 398.4, pp. 1601–1614. DOI: [10.1111/j.1365-2966.2009.14548.x](https://doi.org/10.1111/j.1365-2966.2009.14548.x). arXiv: [0809.3437](https://arxiv.org/abs/0809.3437) [astro-ph].
- Ferrario, Lilia, Andrew Melatos, and Jonathan Zrake (Oct. 2015). “Magnetic Field Generation in Stars”. In: *Space Science Reviews* 191.1-4, pp. 77–109. DOI: [10.1007/s11214-015-0138-y](https://doi.org/10.1007/s11214-015-0138-y). arXiv: [1504.08074](https://arxiv.org/abs/1504.08074) [astro-ph.SR].
- Feynman, R. P. (1957). “An expanded version of the remarks by R. P. Feynman on the reality of gravitational waves”. In: *The Role of Gravitation in Physics: Report from the 1957 Chapel Hill Conference*. Ed. by Dean Rickles and Cécile M. DeWitt. Berlin: Max-Planck-Gesellschaft zur Förderung der Wissenschaften. DOI: [10.34663/9783945561294-34](https://doi.org/10.34663/9783945561294-34).
- Flanagan, Claire S. (May 1990). “Rapid recovery of the Vela pulsar from a giant glitch”. In: *Nature* 345.6274, pp. 416–417. DOI: [10.1038/345416a0](https://doi.org/10.1038/345416a0).
- Ford, John M., Paul Demorest, and Scott Ransom (July 2010). “Heterogeneous real-time computing in radio astronomy”. In: *Software and Cyberinfrastructure for Astronomy*. Ed. by Nicole M. Radziwill and Alan Bridger. Vol. 7740. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 77400A, 77400A. DOI: [10.1117/12.857666](https://doi.org/10.1117/12.857666).

- Fortin, Francis, Federico García, Adolfo Simaz Bunzel, and Sylvain Chaty (Mar. 2023). “A catalogue of high-mass X-ray binaries in the Galaxy: from the INTEGRAL to the Gaia era”. In: *Astronomy and Astrophysics* 671, A149, A149. DOI: [10.1051/0004-6361/202245236](https://doi.org/10.1051/0004-6361/202245236). arXiv: [2302.02656](https://arxiv.org/abs/2302.02656) [astro-ph.HE].
- Franco, Lucia M., Bennett Link, and Richard I. Epstein (Nov. 2000). “Quaking Neutron Stars”. In: *The Astrophysical Journal* 543.2, pp. 987–994. DOI: [10.1086/317121](https://doi.org/10.1086/317121). arXiv: [astro-ph/9911105](https://arxiv.org/abs/astro-ph/9911105) [astro-ph].
- Freire, P. C. C. et al. (Nov. 2011). “Fermi Detection of a Luminous γ -Ray Pulsar in a Globular Cluster”. In: *Science* 334.6059, p. 1107. DOI: [10.1126/science.1207141](https://doi.org/10.1126/science.1207141). arXiv: [1111.3754](https://arxiv.org/abs/1111.3754) [astro-ph.GA].
- Fuentes, J. R., C. M. Espinoza, and A. Reisenegger (Oct. 2019). “Glitch time series and size distributions in eight prolific pulsars”. In: *Astronomy and Astrophysics* 630, A115, A115. DOI: [10.1051/0004-6361/201935939](https://doi.org/10.1051/0004-6361/201935939). arXiv: [1907.09887](https://arxiv.org/abs/1907.09887) [astro-ph.HE].
- Fuentes, J. R., C. M. Espinoza, A. Reisenegger, B. Shaw, B. W. Stappers, and A. G. Lyne (Dec. 2017). “The glitch activity of neutron stars”. In: *Astronomy and Astrophysics* 608, A131, A131. DOI: [10.1051/0004-6361/201731519](https://doi.org/10.1051/0004-6361/201731519). arXiv: [1710.00952](https://arxiv.org/abs/1710.00952) [astro-ph.HE].
- Fujisawa, Kotaro, Shota Kisaka, and Yasufumi Kojima (Nov. 2022). “Magnetically confined mountains on accreting neutron stars with multipole magnetic fields”. In: *Monthly Notices of the Royal Astronomical Society* 516.4, pp. 5196–5208. DOI: [10.1093/mnras/stac2585](https://doi.org/10.1093/mnras/stac2585). arXiv: [2209.04352](https://arxiv.org/abs/2209.04352) [astro-ph.HE].
- Fulgenzi, W., A. Melatos, and B. D. Hughes (Oct. 2017). “Radio pulsar glitches as a state-dependent Poisson process”. In: *Monthly Notices of the Royal Astronomical Society* 470.4, pp. 4307–4329. DOI: [10.1093/mnras/stx1353](https://doi.org/10.1093/mnras/stx1353).
- Gaensler, Bryan M. and Patrick O. Slane (Sept. 2006). “The Evolution and Structure of Pulsar Wind Nebulae”. In: *Annual Review of Astronomy and Astrophysics* 44.1, pp. 17–47. DOI: [10.1146/annurev.astro.44.051905.092528](https://doi.org/10.1146/annurev.astro.44.051905.092528). arXiv: [astro-ph/0601081](https://arxiv.org/abs/astro-ph/0601081) [astro-ph].
- Gancio, G. et al. (Jan. 2020). “Upgraded antennas for pulsar observations in the Argentine Institute of Radio astronomy”. In: *Astronomy and Astrophysics* 633, A84, A84. DOI: [10.1051/0004-6361/201936525](https://doi.org/10.1051/0004-6361/201936525). arXiv: [1908.07049](https://arxiv.org/abs/1908.07049) [astro-ph.IM].
- Gangwar, Yashaswi and David Ian Jones (Aug. 2024). “Applying the starquake model to study the formation of elastic mountains on spinning neutron stars”. In: *Monthly Notices of the Royal Astronomical Society* 532.2, pp. 2763–2777. DOI: [10.1093/mnras/stae1671](https://doi.org/10.1093/mnras/stae1671). arXiv: [2405.00553](https://arxiv.org/abs/2405.00553) [astro-ph.HE].
- Garthwaite, M. C., A. R Hill, S. Yates, and B. Twilley (2013). *Geodetic survey of the Molonglo Radio Observatory*. Tech. rep. Canberra Geoscience Australia. URL: <https://ecat.ga.gov.au/geonetwork/srv/eng/catalog.search#/metadata/75653>.
- Giliberti, E., M. Antonelli, G. Cambiotti, and P. M. Pizzochero (Oct. 2019). “Incompressible analytical models for spinning-down pulsars”. In: *Publications of the Astronomical Society of Australia* 36, e036, e036. DOI: [10.1017/pasa.2019.28](https://doi.org/10.1017/pasa.2019.28). arXiv: [1902.06345](https://arxiv.org/abs/1902.06345) [astro-ph.HE].
- Giliberti, E., G. Cambiotti, M. Antonelli, and P. M. Pizzochero (Jan. 2020). “Modelling strains and stresses in continuously stratified rotating neutron stars”. In: *Monthly Notices of the Royal Astronomical Society* 491.1, pp. 1064–1078. DOI: [10.1093/mnras/stz3099](https://doi.org/10.1093/mnras/stz3099). arXiv: [1809.08542](https://arxiv.org/abs/1809.08542) [astro-ph.HE].

- Gittins, F., N. Andersson, and D. I. Jones (Jan. 2021). “Modelling neutron star mountains”. In: *Monthly Notices of the Royal Astronomical Society* 500.4, pp. 5570–5582. DOI: [10.1093/mnras/staa3635](https://doi.org/10.1093/mnras/staa3635). arXiv: [2009.12794](https://arxiv.org/abs/2009.12794) [astro-ph.HE].
- Glampedakis, Kostas and Nils Andersson (Apr. 2009). “Hydrodynamical Trigger Mechanism for Pulsar Glitches”. In: *Physical Review Letters* 102.14, 141101, p. 141101. DOI: [10.1103/PhysRevLett.102.141101](https://doi.org/10.1103/PhysRevLett.102.141101). arXiv: [0806.3664](https://arxiv.org/abs/0806.3664) [astro-ph].
- Goetz, Evan (2021). *H1 Calibrated Sensitivity Spectra Jan 04 2020 (Representative best of O3b – C01_CLEAN_SUB60HZ)*. LIGO Document LIGO-G2100674. URL: <https://dcc.ligo.org/LIGO-G2100674/public>.
- Goetz, Evan, Ansel Neunzert, Keith Riles, Andrew Matas, Shivaraj Kandhasamy, et al. (2021). *O3 lines and combs found in self-gated C01 data*. LIGO Document T2100200. URL: <https://dcc.ligo.org/T2100200/public>.
- Goetz, Evan and Keith Riles (2023). *Segments used for creating standard SFTs in O3 data*. LIGO Document T2300068. URL: <https://dcc.ligo.org/T2300068/public>.
- Gold, T. (May 1968). “Rotating Neutron Stars as the Origin of the Pulsating Radio Sources”. In: *Nature* 218.5143, pp. 731–732. DOI: [10.1038/218731a0](https://doi.org/10.1038/218731a0).
- Goldreich, Peter and William H. Julian (Aug. 1969). “Pulsar Electrodynamics”. In: *The Astrophysical Journal* 157, p. 869. DOI: [10.1086/150119](https://doi.org/10.1086/150119).
- Goncharov, Boris, Xing-Jiang Zhu, and Eric Thrane (Sept. 2020). “Is there a spectral turnover in the spin noise of millisecond pulsars?” In: *Monthly Notices of the Royal Astronomical Society* 497.3, pp. 3264–3272. DOI: [10.1093/mnras/staa2081](https://doi.org/10.1093/mnras/staa2081). arXiv: [1910.05961](https://arxiv.org/abs/1910.05961) [astro-ph.HE].
- Goncharov, Boris et al. (Mar. 2021). “Identifying and mitigating noise sources in precision pulsar timing data sets”. In: *Monthly Notices of the Royal Astronomical Society* 502.1, pp. 478–493. DOI: [10.1093/mnras/staa3411](https://doi.org/10.1093/mnras/staa3411). arXiv: [2010.06109](https://arxiv.org/abs/2010.06109) [astro-ph.HE].
- Good, D. C. et al. (Nov. 2021). “First Discovery of New Pulsars and RRATs with CHIME/FRB”. in: *The Astrophysical Journal* 922.1, 43, p. 43. DOI: [10.3847/1538-4357/ac1da6](https://doi.org/10.3847/1538-4357/ac1da6). arXiv: [2012.02320](https://arxiv.org/abs/2012.02320) [astro-ph.HE].
- Gravitational Wave Open Science Center (n.d.). *O3 Instrumental Lines*. <https://gwosc.org/O3/o3speclines/>.
- Greenstein, George (Aug. 1970). “Superfluid Turbulence in Neutron Stars”. In: *Nature* 227.5260, pp. 791–794. DOI: [10.1038/227791a0](https://doi.org/10.1038/227791a0).
- Groth, E. J. (Nov. 1975). “Timing of the Crab Pulsar III. The Slowing Down and the Nature of the Random Process”. In: *The Astrophysical Journal Supplement Series* 29, pp. 453–465. DOI: [10.1086/190354](https://doi.org/10.1086/190354).
- Grover, Himanshu, Jaikhomba Singha, Bhal Chandra Joshi, and Paramasivan Arumugam (Sept. 2022). “Detection of a glitch event in PSR J0742-2822 observed using the Automated Glitch Detection Pipeline implemented at the Ooty Radio Telescope (ORT)”. in: *The Astronomer’s Telegram* 15629, p. 1.
- Grover, Himanshu et al. (Jan. 2023). “Confirmation of glitch event detected in PSR J1740-3015 using the upgraded Giant Metrewave Radio Telescope (uGMRT)”. in: *The Astronomer’s Telegram* 15851, p. 1.
- Gügercinoglu, E., M. Y. Ge, J. P. Yuan, and S. Q. Zhou (Mar. 2022). “Glitches in four gamma-ray pulsars and inferences on the neutron star structure”. In: *Monthly Notices of the Royal Astronomical Society* 511.1, pp. 425–439. DOI: [10.1093/mnras/stac026](https://doi.org/10.1093/mnras/stac026).

- Gullahorn, G. E. and J. M. Rankin (Sept. 1982). “Pulsar timing irregularities.” In: *The Astrophysical Journal* 260, pp. 520–537. DOI: [10.1086/160274](https://doi.org/10.1086/160274).
- Haberl, Frank (Apr. 2007). “The magnificent seven: magnetic fields and surface temperature distributions”. In: *Astrophysics and Space Science* 308.1-4, pp. 181–190. DOI: [10.1007/s10509-007-9342-x](https://doi.org/10.1007/s10509-007-9342-x). arXiv: [astro-ph/0609066](https://arxiv.org/abs/astro-ph/0609066) [astro-ph].
- Hankins, T. H. and B. J. Rickett (Jan. 1975). “Pulsar signal processing.” In: *Methods in Computational Physics* 14, pp. 55–129. DOI: [10.1016/B978-0-12-460814-6.50007-3](https://doi.org/10.1016/B978-0-12-460814-6.50007-3).
- Harris, William E. (Oct. 1996). “A Catalog of Parameters for Globular Clusters in the Milky Way”. In: *The Astronomical Journal* 112, p. 1487. DOI: [10.1086/118116](https://doi.org/10.1086/118116).
- Haskell, B., D. Antonopoulou, and C. Barenghi (Nov. 2020). “Turbulent, pinned superfluids in neutron stars and pulsar glitch recoveries”. In: *Monthly Notices of the Royal Astronomical Society* 499.1, pp. 161–170. DOI: [10.1093/mnras/staa2678](https://doi.org/10.1093/mnras/staa2678). arXiv: [2007.02748](https://arxiv.org/abs/2007.02748) [astro-ph.HE].
- Haskell, B., P. M. Pizzochero, and T. Sidery (Feb. 2012). “Modelling pulsar glitches with realistic pinning forces: a hydrodynamical approach”. In: *Monthly Notices of the Royal Astronomical Society* 420.1, pp. 658–671. DOI: [10.1111/j.1365-2966.2011.20080.x](https://doi.org/10.1111/j.1365-2966.2011.20080.x). arXiv: [1107.5295](https://arxiv.org/abs/1107.5295) [astro-ph.SR].
- Haskell, Brynmor and Andrew Melatos (Jan. 2015). “Models of pulsar glitches”. In: *International Journal of Modern Physics D* 24.3, 1530008, p. 1530008. DOI: [10.1142/S0218271815300086](https://doi.org/10.1142/S0218271815300086). arXiv: [1502.07062](https://arxiv.org/abs/1502.07062) [astro-ph.SR].
- Haslam, C. G. T., C. J. Salter, H. Stoffel, and W. E. Wilson (Jan. 1982). “A 408-MHZ All-Sky Continuum Survey. II. The Atlas of Contour Maps”. In: *Astronomy and Astrophysics Supplement Series* 47, p. 1.
- Hayes, Jeremiah F (2002). “The Viterbi algorithm applied to digital data transmission”. In: *IEEE Communications Magazine* 40.5, pp. 26–32.
- Helfand, D. J., R. N. Manchester, and J. H. Taylor (June 1975). “Observations of pulsar radio emission. III. Stability of integrated profiles.” In: *The Astrophysical Journal* 198, pp. 661–670. DOI: [10.1086/153644](https://doi.org/10.1086/153644).
- H.E.S.S. Collaboration et al. (Nov. 2023). “Discovery of a radiation component from the Vela pulsar reaching 20 teraelectronvolts”. In: *Nature Astronomy* 7, pp. 1341–1350. DOI: [10.1038/s41550-023-02052-3](https://doi.org/10.1038/s41550-023-02052-3). arXiv: [2310.06181](https://arxiv.org/abs/2310.06181) [astro-ph.HE].
- Hessels, Jason W. T., Scott M. Ransom, Ingrid H. Stairs, Paulo C. C. Freire, Victoria M. Kaspi, and Fernando Camilo (Mar. 2006). “A Radio Pulsar Spinning at 716 Hz”. In: *Science* 311.5769, pp. 1901–1904. DOI: [10.1126/science.1123430](https://doi.org/10.1126/science.1123430). arXiv: [astro-ph/0601337](https://arxiv.org/abs/astro-ph/0601337) [astro-ph].
- Hewish, A., S. J. Bell, J. D. H. Pilkington, P. F. Scott, and R. A. Collins (Feb. 1968). “Observation of a Rapidly Pulsating Radio Source”. In: *Nature* 217.5130, pp. 709–713. DOI: [10.1038/217709a0](https://doi.org/10.1038/217709a0).
- Ho, W. C. G., C. M. Espinoza, D. Antonopoulou, and N. Andersson (Oct. 2015). “Pinning down the superfluid and measuring masses using pulsar glitches”. In: *Science Advances* 1.9, e1500578–e1500578. DOI: [10.1126/sciadv.1500578](https://doi.org/10.1126/sciadv.1500578). arXiv: [1510.00395](https://arxiv.org/abs/1510.00395) [astro-ph.SR].
- Ho, Wynn C. G. et al. (Nov. 2020). “Return of the Big Glitcher: NICER timing and glitches of PSR J0537-6910”. In: *Monthly Notices of the Royal Astronomical Society* 498.4, pp. 4605–4614. DOI: [10.1093/mnras/staa2640](https://doi.org/10.1093/mnras/staa2640). arXiv: [2009.00030](https://arxiv.org/abs/2009.00030) [astro-ph.HE].

- Hobbs, G., A. G. Lyne, and M. Kramer (Feb. 2010). “An analysis of the timing irregularities for 366 pulsars”. In: *Monthly Notices of the Royal Astronomical Society* 402.2, pp. 1027–1048. DOI: [10.1111/j.1365-2966.2009.15938.x](https://doi.org/10.1111/j.1365-2966.2009.15938.x). arXiv: [0912.4537](https://arxiv.org/abs/0912.4537) [astro-ph.GA].
- Hobbs, G., A. G. Lyne, M. Kramer, C. E. Martin, and C. Jordan (Oct. 2004). “Long-term timing observations of 374 pulsars”. In: *Monthly Notices of the Royal Astronomical Society* 353.4, pp. 1311–1344. DOI: [10.1111/j.1365-2966.2004.08157.x](https://doi.org/10.1111/j.1365-2966.2004.08157.x).
- Hobbs, G. et al. (Aug. 2011). “The Parkes Observatory Pulsar Data Archive”. In: *Publications of the Astronomical Society of Australia* 28.3, pp. 202–214. DOI: [10.1071/AS11016](https://doi.org/10.1071/AS11016). arXiv: [1105.5746](https://arxiv.org/abs/1105.5746) [astro-ph.IM].
- Hobbs, G. B., R. T. Edwards, and R. N. Manchester (June 2006). “TEMPO2, a new pulsar-timing package - I. An overview”. In: *Monthly Notices of the Royal Astronomical Society* 369.2, pp. 655–672. DOI: [10.1111/j.1365-2966.2006.10302.x](https://doi.org/10.1111/j.1365-2966.2006.10302.x). arXiv: [astro-ph/0603381](https://arxiv.org/abs/astro-ph/0603381) [astro-ph].
- Hobbs, George, Rui Luo, and Lawrence Toomey (May 2022). *simulateSearch: High-time resolution data sets simulations for radio telescopes*. Astrophysics Source Code Library, record ascl:2205.025.
- Hobbs, George et al. (Apr. 2020). “An ultra-wide bandwidth (704 to 4 032 MHz) receiver for the Parkes radio telescope”. In: *Publications of the Astronomical Society of Australia* 37, e012, e012. DOI: [10.1017/pasa.2020.2](https://doi.org/10.1017/pasa.2020.2). arXiv: [1911.00656](https://arxiv.org/abs/1911.00656) [astro-ph.IM].
- Horowitz, C. J. and Kai Kadau (May 2009). “Breaking Strain of Neutron Star Crust and Gravitational Waves”. In: *Physical Review Letters* 102.19, 191102, p. 191102. DOI: [10.1103/PhysRevLett.102.191102](https://doi.org/10.1103/PhysRevLett.102.191102). arXiv: [0904.1986](https://arxiv.org/abs/0904.1986) [astro-ph.SR].
- Hotan, A. W., W. van Straten, and R. N. Manchester (Jan. 2004). “PSRCHIVE and PSRFITS: An Open Approach to Radio Pulsar Data Storage and Analysis”. In: *Publications of the Astronomical Society of Australia* 21.3, pp. 302–309. DOI: [10.1071/AS04022](https://doi.org/10.1071/AS04022). arXiv: [astro-ph/0404549](https://arxiv.org/abs/astro-ph/0404549) [astro-ph].
- Howitt, G., A. Melatos, and A. Delaigle (Nov. 2018). “Nonparametric Estimation of the Size and Waiting Time Distributions of Pulsar Glitches”. In: *The Astrophysical Journal* 867.1, 60, p. 60. DOI: [10.3847/1538-4357/aae20a](https://doi.org/10.3847/1538-4357/aae20a). arXiv: [1809.06038](https://arxiv.org/abs/1809.06038) [astro-ph.HE].
- Howitt, G., A. Melatos, and B. Haskell (Oct. 2020). “Simulating pulsar glitches: an N-body solver for superfluid vortex motion in two dimensions”. In: *Monthly Notices of the Royal Astronomical Society* 498.1, pp. 320–331. DOI: [10.1093/mnras/staa2314](https://doi.org/10.1093/mnras/staa2314). arXiv: [2008.00365](https://arxiv.org/abs/2008.00365) [astro-ph.HE].
- Hulse, R. A. and J. H. Taylor (Jan. 1975). “Discovery of a pulsar in a binary system.” In: *The Astrophysical Journal* 195, pp. L51–L53. DOI: [10.1086/181708](https://doi.org/10.1086/181708).
- Hurley, K. et al. (Apr. 2005). “An exceptionally bright flare from SGR 1806-20 and the origins of short-duration γ -ray bursts”. In: *Nature* 434.7037, pp. 1098–1103. DOI: [10.1038/nature03519](https://doi.org/10.1038/nature03519). arXiv: [astro-ph/0502329](https://arxiv.org/abs/astro-ph/0502329) [astro-ph].
- Hutchins, T. J. and D. I. Jones (June 2023). “Gravitational radiation from thermal mountains on accreting neutron stars: sources of temperature non-axisymmetry”. In: *Monthly Notices of the Royal Astronomical Society* 522.1, pp. 226–251. DOI: [10.1093/mnras/stad967](https://doi.org/10.1093/mnras/stad967). arXiv: [2212.07452](https://arxiv.org/abs/2212.07452) [astro-ph.HE].
- Ioka, K. (Oct. 2001). “Magnetic deformation of magnetars for the giant flares of the soft gamma-ray repeaters”. In: *Monthly Notices of the Royal Astronomical Society* 327.2, pp. 639–662. DOI: [10.1046/j.1365-8711.2001.04756.x](https://doi.org/10.1046/j.1365-8711.2001.04756.x). arXiv: [astro-ph/0009327](https://arxiv.org/abs/astro-ph/0009327) [astro-ph].

- Ivanova, N., C. O. Heinke, F. A. Rasio, K. Belczynski, and J. M. Fregeau (May 2008). “Formation and evolution of compact binaries in globular clusters - II. Binaries with neutron stars”. In: *Monthly Notices of the Royal Astronomical Society* 386.1, pp. 553–576. DOI: [10.1111/j.1365-2966.2008.13064.x](https://doi.org/10.1111/j.1365-2966.2008.13064.x). arXiv: [0706.4096](https://arxiv.org/abs/0706.4096) [astro-ph].
- Jankowski, F., E. F. Keane, and B. W. Stappers (June 2021). “Constraints on wide-band radiative changes after a glitch in PSR J1452-6036”. In: *Monthly Notices of the Royal Astronomical Society* 504.1, pp. 406–415. DOI: [10.1093/mnras/stab824](https://doi.org/10.1093/mnras/stab824). arXiv: [2103.09869](https://arxiv.org/abs/2103.09869) [astro-ph.HE].
- Jankowski, F., W. van Straten, E. F. Keane, M. Bailes, E. D. Barr, S. Johnston, and M. Kerr (Feb. 2018). “Spectral properties of 441 radio pulsars”. In: *Monthly Notices of the Royal Astronomical Society* 473.4, pp. 4436–4458. DOI: [10.1093/mnras/stx2476](https://doi.org/10.1093/mnras/stx2476). arXiv: [1709.08864](https://arxiv.org/abs/1709.08864) [astro-ph.HE].
- Jankowski, F. et al. (Apr. 2019). “The UTMOST pulsar timing programme I: Overview and first results”. In: *Monthly Notices of the Royal Astronomical Society* 484.3, pp. 3691–3712. DOI: [10.1093/mnras/sty3390](https://doi.org/10.1093/mnras/sty3390). arXiv: [1812.04038](https://arxiv.org/abs/1812.04038) [astro-ph.HE].
- Jankowski, Fabian et al. (Nov. 2015a). “Glitch event observed in the pulsar PSR J1740-3015”. In: *The Astronomer’s Telegram* 8298, p. 1.
- Jankowski, Fabian et al. (Jan. 2015b). “Glitch event observed in the Vela pulsar (PSR J0835-4510)”. In: *The Astronomer’s Telegram* 6903, p. 1.
- Jankowski, Fabian et al. (May 2016). “Glitch event observed in the pulsar PSR J1740-3015”. In: *The Astronomer’s Telegram* 9054, p. 1.
- Jankowski, Fabian et al. (Sept. 2017). “Large glitch event observed in PSR J1731-4744”. In: *The Astronomer’s Telegram* 10770, p. 1.
- Janssen, G. H. and B. W. Stappers (Oct. 2006). “30 glitches in slow pulsars”. In: *Astronomy and Astrophysics* 457.2, pp. 611–618. DOI: [10.1051/0004-6361:20065267](https://doi.org/10.1051/0004-6361:20065267). arXiv: [astro-ph/0607260](https://arxiv.org/abs/astro-ph/0607260) [astro-ph].
- Jaranowski, Piotr, Andrzej Królak, and Bernard F. Schutz (Sept. 1998). “Data analysis of gravitational-wave signals from spinning neutron stars: The signal and its detection”. In: *Physical Review D* 58.6, 063001, p. 063001. DOI: [10.1103/PhysRevD.58.063001](https://doi.org/10.1103/PhysRevD.58.063001). arXiv: [gr-qc/9804014](https://arxiv.org/abs/gr-qc/9804014) [gr-qc].
- Jaume, Rafel, Rodrigo Tenorio, and Alicia M. Sintes (Mar. 2024). “Assessing the Similarity of Continuous Gravitational-Wave Signals to Narrow Instrumental Artifacts”. In: *Universe* 10.3, 121, p. 121. DOI: [10.3390/universe10030121](https://doi.org/10.3390/universe10030121). arXiv: [2403.03027](https://arxiv.org/abs/2403.03027) [gr-qc].
- Jennings, Ross J., James M. Cordes, and Shami Chatterjee (Dec. 2020). “Pulsar Timing Signatures of Circumbinary Asteroid Belts”. In: *The Astrophysical Journal* 904.2, 191, p. 191. DOI: [10.3847/1538-4357/abc178](https://doi.org/10.3847/1538-4357/abc178). arXiv: [2007.10388](https://arxiv.org/abs/2007.10388) [astro-ph.HE].
- Jennings, Ross J., Nanograv Collaboration, and Chime/Pulsar Collaboration (Feb. 2022). “Recovery of PSR J1713+0747 from a sustained pulse shape change”. In: *The Astronomer’s Telegram* 15223, p. 1.
- Johnson, Aaron D. et al. (May 2024). “NANOGrav 15-year gravitational-wave background methods”. In: *Physical Review D* 109.10, 103012, p. 103012. DOI: [10.1103/PhysRevD.109.103012](https://doi.org/10.1103/PhysRevD.109.103012). arXiv: [2306.16223](https://arxiv.org/abs/2306.16223) [astro-ph.HE].
- Johnson-McDaniel, Nathan K. and Benjamin J. Owen (Aug. 2013). “Maximum elastic deformations of relativistic stars”. In: *Physical Review D* 88.4, 044004, p. 044004. DOI: [10.1103/PhysRevD.88.044004](https://doi.org/10.1103/PhysRevD.88.044004). arXiv: [1208.5227](https://arxiv.org/abs/1208.5227) [astro-ph.SR].

- Johnston, Simon et al. (Apr. 2020). “The Thousand-Pulsar-Array programme on MeerKAT - I. Science objectives and first results”. In: *Monthly Notices of the Royal Astronomical Society* 493.3, pp. 3608–3615. DOI: [10.1093/mnras/staa516](https://doi.org/10.1093/mnras/staa516). arXiv: [2002.10250](https://arxiv.org/abs/2002.10250) [astro-ph.HE].
- Jones, D. I. (Aug. 2004). “Is timing noise important in the gravitational wave detection of neutron stars?” In: *Physical Review D* 70.4, 042002, p. 042002. DOI: [10.1103/PhysRevD.70.042002](https://doi.org/10.1103/PhysRevD.70.042002). arXiv: [gr-qc/0406045](https://arxiv.org/abs/gr-qc/0406045) [gr-qc].
- Jones, Dana and Ling Sun (Jan. 2021). “Search for continuous gravitational waves from Fomalhaut b in the second Advanced LIGO observing run with a hidden Markov model”. In: *Physical Review D* 103.2, 023020, p. 023020. DOI: [10.1103/PhysRevD.103.023020](https://doi.org/10.1103/PhysRevD.103.023020). arXiv: [2007.08732](https://arxiv.org/abs/2007.08732) [gr-qc].
- Jones, Dana et al. (Dec. 2022). “Validating continuous gravitational-wave candidates from a semicoherent search using Doppler modulation and an effective point spread function”. In: *Physical Review D* 106.12, 123011, p. 123011. DOI: [10.1103/PhysRevD.106.123011](https://doi.org/10.1103/PhysRevD.106.123011). arXiv: [2203.14468](https://arxiv.org/abs/2203.14468) [gr-qc].
- Jones, P. B. (Sept. 1990a). “Rotation of the neutron-drip superfluid in pulsars : period discontinuities and internal temperatures.” In: *Monthly Notices of the Royal Astronomical Society* 246, p. 315.
- (Sept. 1990b). “The generation of timing noise by superfluid rotation in pulsars.” In: *Monthly Notices of the Royal Astronomical Society* 246, p. 364.
- (May 1998). “The origin of pulsar glitches”. In: *Monthly Notices of the Royal Astronomical Society* 296.1, pp. 217–224. DOI: [10.1046/j.1365-8711.1998.01464.x](https://doi.org/10.1046/j.1365-8711.1998.01464.x).
- Kagra Collaboration et al. (Jan. 2019). “KAGRA: 2.5 generation interferometric gravitational wave detector”. In: *Nature Astronomy* 3, pp. 35–40. DOI: [10.1038/s41550-018-0658-y](https://doi.org/10.1038/s41550-018-0658-y). arXiv: [1811.08079](https://arxiv.org/abs/1811.08079) [gr-qc].
- Kalman, Rudolph Emil (1960). “A new approach to linear filtering and prediction problems”. In: *Transactions of the ASME – Journal of Basic Engineering* 82.1, pp. 35–45.
- Kaspi, Victoria M. and Andrei M. Beloborodov (Aug. 2017). “Magnetars”. In: *Annual Review of Astronomy and Astrophysics* 55.1, pp. 261–301. DOI: [10.1146/annurev-astro-081915-023329](https://doi.org/10.1146/annurev-astro-081915-023329). arXiv: [1703.00068](https://arxiv.org/abs/1703.00068) [astro-ph.HE].
- Katz, J. I. (Feb. 1975). “Two kinds of stellar collapse”. In: *Nature* 253.5494, pp. 698–699. DOI: [10.1038/253698a0](https://doi.org/10.1038/253698a0).
- Keane, E. F. (Nov. 2010). “The Transient Radio Sky”. PhD thesis. University of Manchester, UK.
- Keith, M. J. et al. (May 2024). “The Thousand-Pulsar-Array programme on MeerKAT - XIII. Timing, flux density, rotation measure, and dispersion measure time series of 597 pulsars”. In: *Monthly Notices of the Royal Astronomical Society* 530.2, pp. 1581–1591. DOI: [10.1093/mnras/stae937](https://doi.org/10.1093/mnras/stae937). arXiv: [2404.02051](https://arxiv.org/abs/2404.02051) [astro-ph.HE].
- Keith, Michael J. and Iuliana C. Nițu (Aug. 2023). “Impact of quasi-periodic and steep-spectrum timing noise on the measurement of pulsar timing parameters”. In: *Monthly Notices of the Royal Astronomical Society* 523.3, pp. 4603–4614. DOI: [10.1093/mnras/stad1713](https://doi.org/10.1093/mnras/stad1713). arXiv: [2306.03529](https://arxiv.org/abs/2306.03529) [astro-ph.HE].
- Kerin, A. D. and A. Melatos (Aug. 2022). “Mountain formation by repeated, inhomogeneous crustal failure in a neutron star”. In: *Monthly Notices of the Royal Astronomical Society* 514.2, pp. 1628–1644. DOI: [10.1093/mnras/stac1351](https://doi.org/10.1093/mnras/stac1351). arXiv: [2205.15026](https://arxiv.org/abs/2205.15026) [astro-ph.HE].

- Kerr, M. (Feb. 2019). “Fermi LAT Detection of the Recent Glitch in the Vela Pulsar (PSR J0835-4510)”. In: *The Astronomer’s Telegram* 12481, p. 1.
- Kerr, M., S. Johnston, G. Hobbs, and R. M. Shannon (Aug. 2015). “Limits on Planet Formation Around Young Pulsars and Implications for Supernova fallback Disks”. In: *The Astrophysical Journal* 809.1, L11, p. L11. DOI: [10.1088/2041-8205/809/1/L11](https://doi.org/10.1088/2041-8205/809/1/L11). arXiv: [1507.06982](https://arxiv.org/abs/1507.06982) [astro-ph.HE].
- Khomenko, V., M. Antonelli, and B. Haskell (Dec. 2019). “Hydrodynamical instabilities in the superfluid interior of neutron stars with background flows between the components”. In: *Physical Review D* 100.12, 123002, p. 123002. DOI: [10.1103/PhysRevD.100.123002](https://doi.org/10.1103/PhysRevD.100.123002). arXiv: [1906.05177](https://arxiv.org/abs/1906.05177) [astro-ph.HE].
- Khomenko, V. and B. Haskell (May 2018). “Modelling Pulsar Glitches: The Hydrodynamics of Superfluid Vortex Avalanches in Neutron Stars”. In: *Publications of the Astronomical Society of Australia* 35, e020, e020. DOI: [10.1017/pasa.2018.12](https://doi.org/10.1017/pasa.2018.12). arXiv: [1801.01413](https://arxiv.org/abs/1801.01413) [astro-ph.HE].
- Kimpson, T. et al. (2024). “Adaptive cancellation of mains power interference in continuous gravitational wave searches with a hidden Markov model”. In: *submitted*.
- Knee, Alan M. et al. (Mar. 2024). “Search for continuous gravitational waves directed at subthreshold radiometer candidates in O3 LIGO data”. In: *Physical Review D* 109.6, 062008, p. 062008. DOI: [10.1103/PhysRevD.109.062008](https://doi.org/10.1103/PhysRevD.109.062008). arXiv: [2311.12138](https://arxiv.org/abs/2311.12138) [gr-qc].
- Knight, H. S., M. Bailes, R. N. Manchester, S. M. Ord, and B. A. Jacoby (Apr. 2006). “Green Bank Telescope Studies of Giant Pulses from Millisecond Pulsars”. In: *The Astrophysical Journal* 640.2, pp. 941–949. DOI: [10.1086/500292](https://doi.org/10.1086/500292). arXiv: [astro-ph/0512341](https://arxiv.org/abs/astro-ph/0512341) [astro-ph].
- Konar, Sushan and Dipankar Bhattacharya (Jan. 1997). “Magnetic field evolution of accreting neutron stars”. In: *Monthly Notices of the Royal Astronomical Society* 284.2, pp. 311–317. DOI: [10.1093/mnras/284.2.311](https://doi.org/10.1093/mnras/284.2.311).
- Kouveliotou, C., T. Strohmayer, K. Hurley, J. van Paradijs, M. H. Finger, S. Dieters, P. Woods, C. Thompson, and R. C. Duncan (Jan. 1999). “Discovery of a Magnetar Associated with the Soft Gamma Repeater SGR 1900+14”. In: *The Astrophysical Journal* 510.2, pp. L115–L118. DOI: [10.1086/311813](https://doi.org/10.1086/311813). arXiv: [astro-ph/9809140](https://arxiv.org/abs/astro-ph/9809140) [astro-ph].
- Kremer, Kyle, Jim Fuller, Anthony L. Piro, and Scott M. Ransom (Oct. 2023). “Connecting the young pulsars in Milky Way globular clusters with white dwarf mergers and the M81 fast radio burst”. In: *Monthly Notices of the Royal Astronomical Society* 525.1, pp. L22–L27. DOI: [10.1093/mnras/slad088](https://doi.org/10.1093/mnras/slad088). arXiv: [2305.11933](https://arxiv.org/abs/2305.11933) [astro-ph.HE].
- Krishnan, Badri, Alicia M. Sintes, Maria Alessandra Papa, Bernard F. Schutz, Sergio Frasca, and Cristiano Palomba (Oct. 2004). “Hough transform search for continuous gravitational waves”. In: *Physical Review D* 70.8, 082001, p. 082001. DOI: [10.1103/PhysRevD.70.082001](https://doi.org/10.1103/PhysRevD.70.082001). arXiv: [gr-qc/0407001](https://arxiv.org/abs/gr-qc/0407001) [gr-qc].
- Kulkarni, S. R. (July 2020). “Dispersion measure: Confusion, Constants & Clarity”. In: *arXiv e-prints*, arXiv:2007.02886, arXiv:2007.02886. DOI: [10.48550/arXiv.2007.02886](https://doi.org/10.48550/arXiv.2007.02886). arXiv: [2007.02886](https://arxiv.org/abs/2007.02886) [astro-ph.HE].
- Kumamoto, H. et al. (Mar. 2021). “Flux density variability of 286 radio pulsars from a decade of monitoring”. In: *Monthly Notices of the Royal Astronomical Society* 501.3, pp. 4490–4513. DOI: [10.1093/mnras/staa3910](https://doi.org/10.1093/mnras/staa3910). arXiv: [2012.08119](https://arxiv.org/abs/2012.08119) [astro-ph.HE].
- Large, M. I., A. E. Vaughan, and B. Y. Mills (Oct. 1968). “A Pulsar Supernova Association?” In: *Nature* 220.5165, pp. 340–341. DOI: [10.1038/220340a0](https://doi.org/10.1038/220340a0).

- Large, M. L. and A. E. Vaughan (Jan. 1971). “A search of the galactic plane for high dispersion pulsars”. In: *Monthly Notices of the Royal Astronomical Society* 151, p. 277. DOI: [10.1093/mnras/151.3.277](https://doi.org/10.1093/mnras/151.3.277).
- Lasky, Paul D. and Andrew Melatos (Nov. 2013). “Tilted torus magnetic fields in neutron stars and their gravitational wave signatures”. In: *Physical Review D* 88.10, 103005, p. 103005. DOI: [10.1103/PhysRevD.88.103005](https://doi.org/10.1103/PhysRevD.88.103005). arXiv: [1310.7633](https://arxiv.org/abs/1310.7633) [astro-ph.HE].
- Leaci, Paola and Reinhard Prix (May 2015). “Directed searches for continuous gravitational waves from binary systems: Parameter-space metrics and optimal Scorpius X-1 sensitivity”. In: *Physical Review D* 91.10, 102003, p. 102003. DOI: [10.1103/PhysRevD.91.102003](https://doi.org/10.1103/PhysRevD.91.102003). arXiv: [1502.00914](https://arxiv.org/abs/1502.00914) [gr-qc].
- Lee, Jongsu, C. Y. Hui, J. Takata, A. K. H. Kong, Pak-Hin Thomas Tam, Kwan-Lok Li, and K. S. Cheng (Feb. 2023). “A Comparison of Millisecond Pulsar Populations between Globular Clusters and the Galactic Field”. In: *The Astrophysical Journal* 944.2, 225, p. 225. DOI: [10.3847/1538-4357/acb5a3](https://doi.org/10.3847/1538-4357/acb5a3). arXiv: [2302.08776](https://arxiv.org/abs/2302.08776) [astro-ph.HE].
- Lentati, L., P. Alexander, M. P. Hobson, F. Feroz, R. van Haasteren, K. J. Lee, and R. M. Shannon (Jan. 2014). “TEMPONEST: a Bayesian approach to pulsar timing analysis”. In: *Monthly Notices of the Royal Astronomical Society* 437.3, pp. 3004–3023. DOI: [10.1093/mnras/stt2122](https://doi.org/10.1093/mnras/stt2122). arXiv: [1310.2120](https://arxiv.org/abs/1310.2120) [astro-ph.IM].
- Lentati, Lindley, P. Alexander, M. P. Hobson, S. Taylor, J. Gair, S. T. Balan, and R. van Haasteren (May 2013). “Hyper-efficient model-independent Bayesian method for the analysis of pulsar timing data”. In: *Physical Review D* 87.10, 104021, p. 104021. DOI: [10.1103/PhysRevD.87.104021](https://doi.org/10.1103/PhysRevD.87.104021). arXiv: [1210.3578](https://arxiv.org/abs/1210.3578) [astro-ph.IM].
- Levin, L., A. G. Lyne, G. Desvignes, R. P. Eatough, R. Karuppusamy, M. Kramer, M. Mickaliger, B. W. Stappers, and P. Weltevrede (Oct. 2019). “Spin frequency evolution and pulse profile variations of the recently re-activated radio magnetar XTE J1810-197”. In: *Monthly Notices of the Royal Astronomical Society* 488.4, pp. 5251–5258. DOI: [10.1093/mnras/stz2074](https://doi.org/10.1093/mnras/stz2074). arXiv: [1903.02660](https://arxiv.org/abs/1903.02660) [astro-ph.HE].
- Levin, Lina et al. (Sept. 2010). “A Radio-loud Magnetar in X-ray Quiescence”. In: *The Astrophysical Journal* 721.1, pp. L33–L37. DOI: [10.1088/2041-8205/721/1/L33](https://doi.org/10.1088/2041-8205/721/1/L33). arXiv: [1007.1052](https://arxiv.org/abs/1007.1052) [astro-ph.HE].
- Levy, Bernard C (2008). *Principles of signal detection and parameter estimation*. Springer Science & Business Media.
- Li, Wei et al. (Oct. 2023). “Results of 23 yr of Pulsar Timing of PSR J1453-6413”. In: *Research in Astronomy and Astrophysics* 23.10, 105014, p. 105014. DOI: [10.1088/1674-4527/acf1e1](https://doi.org/10.1088/1674-4527/acf1e1). arXiv: [2306.13611](https://arxiv.org/abs/2306.13611) [astro-ph.HE].
- LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration (2018). *LVK Algorithm Library - LALSuite*. Free software (GPL). DOI: [10.7935/GT1W-FZ16](https://doi.org/10.7935/GT1W-FZ16).
- Link, B. (Nov. 2006). “Incompatibility of long-period neutron star precession with creeping neutron vortices”. In: *Astronomy and Astrophysics* 458.3, pp. 881–884. DOI: [10.1051/0004-6361:20065664](https://doi.org/10.1051/0004-6361:20065664). arXiv: [astro-ph/0608319](https://arxiv.org/abs/astro-ph/0608319) [astro-ph].
- (May 2012). “Instability of superfluid flow in the neutron star inner crust”. In: *Monthly Notices of the Royal Astronomical Society* 422.2, pp. 1640–1647. DOI: [10.1111/j.1365-2966.2012.20740.x](https://doi.org/10.1111/j.1365-2966.2012.20740.x). arXiv: [1105.4654](https://arxiv.org/abs/1105.4654) [astro-ph.SR].
- Link, Bennett (Sept. 2003). “Constraining Hadronic Superfluidity with Neutron Star Precession”. In: *Physical Review Letters* 91.10, 101101, p. 101101. DOI: [10.1103/PhysRevLett.91.101101](https://doi.org/10.1103/PhysRevLett.91.101101). arXiv: [astro-ph/0302441](https://arxiv.org/abs/astro-ph/0302441) [astro-ph].

- Link, Bennett (July 2014). “Thermally Activated Post-glitch Response of the Neutron Star Inner Crust and Core. I. Theory”. In: *The Astrophysical Journal* 789.2, 141, p. 141. DOI: [10.1088/0004-637X/789/2/141](https://doi.org/10.1088/0004-637X/789/2/141). arXiv: [1311.2499](https://arxiv.org/abs/1311.2499) [astro-ph.SR].
- Link, Bennett and Richard I. Epstein (Feb. 1996). “Thermally Driven Neutron Star Glitches”. In: *The Astrophysical Journal* 457, p. 844. DOI: [10.1086/176779](https://doi.org/10.1086/176779). arXiv: [astro-ph/9508021](https://arxiv.org/abs/astro-ph/9508021) [astro-ph].
- Link, Bennett, Richard I. Epstein, and Gordon Baym (Jan. 1993). “Superfluid Vortex Creep and Rotational Dynamics of Neutron Stars”. In: *The Astrophysical Journal* 403, p. 285. DOI: [10.1086/172202](https://doi.org/10.1086/172202).
- Link, Bennett K. and Richard I. Epstein (June 1991). “Mechanics and Energetics of Vortex Unpinning in Neutron Stars”. In: *The Astrophysical Journal* 373, p. 592. DOI: [10.1086/170078](https://doi.org/10.1086/170078).
- Liu, K., J. P. W. Verbiest, M. Kramer, B. W. Stappers, W. van Straten, and J. M. Cordes (Nov. 2011). “Prospects for high-precision pulsar timing”. In: *Monthly Notices of the Royal Astronomical Society* 417.4, pp. 2916–2926. DOI: [10.1111/j.1365-2966.2011.19452.x](https://doi.org/10.1111/j.1365-2966.2011.19452.x). arXiv: [1107.3086](https://arxiv.org/abs/1107.3086) [astro-ph.HE].
- Liu, P. et al. (Dec. 2023). “Pulse profile variability associated with the glitch of PSR J1048–5832”. In: *arXiv e-prints*, arXiv:2312.04305, arXiv:2312.04305. DOI: [10.48550/arXiv.2312.04305](https://doi.org/10.48550/arXiv.2312.04305). arXiv: [2312.04305](https://arxiv.org/abs/2312.04305) [astro-ph.HE].
- Liu, Y., M. J. Keith, D. Antonopoulou, P. Weltevrede, B. Shaw, B. W. Stappers, A. G. Lyne, M. B. Mickaliger, and A. Basu (July 2024). “Measuring glitch recoveries and braking indices with Bayesian model selection”. In: *Monthly Notices of the Royal Astronomical Society* 532.1, pp. 859–882. DOI: [10.1093/mnras/stae1499](https://doi.org/10.1093/mnras/stae1499). arXiv: [2406.09219](https://arxiv.org/abs/2406.09219) [astro-ph.HE].
- Lopez Armengol, F. G., C. O. Lousto, S. del Palacio, F. Garcia, L. Combi, J. A. Combi, G. Gancio, A. L. Mueller, and P. Kornecki (Feb. 2019). “Radio observations following the recent glitch of Vela Pulsar (PSR B0833-45)”. In: *The Astronomer’s Telegram* 12482, p. 1.
- Lorimer, D. R. and M. Kramer (2004). *Handbook of Pulsar Astronomy*. Vol. 4.
- Lower, M. E. et al. (May 2020). “The UTMOST pulsar timing programme - II. Timing noise across the pulsar population”. In: *Monthly Notices of the Royal Astronomical Society* 494.1, pp. 228–245. DOI: [10.1093/mnras/staa615](https://doi.org/10.1093/mnras/staa615). arXiv: [2002.12481](https://arxiv.org/abs/2002.12481) [astro-ph.HE].
- Lower, M. E. et al. (Dec. 2021). “The impact of glitches on young pulsar rotational evolution”. In: *Monthly Notices of the Royal Astronomical Society* 508.3, pp. 3251–3274. DOI: [10.1093/mnras/stab2678](https://doi.org/10.1093/mnras/stab2678). arXiv: [2109.07612](https://arxiv.org/abs/2109.07612) [astro-ph.HE].
- Lower, M. E. et al. (Oct. 2023a). “Rotational and radio emission properties of PSR J0738-4042 over half a century”. In: *Monthly Notices of the Royal Astronomical Society* 524.4, pp. 5904–5917. DOI: [10.1093/mnras/stad2243](https://doi.org/10.1093/mnras/stad2243). arXiv: [2307.11953](https://arxiv.org/abs/2307.11953) [astro-ph.HE].
- Lower, Marcus E. (2022). “Exploring the magnetospheric and rotational properties of radio pulsars”. PhD thesis. Swinburne University of Technology.
- Lower, Marcus E. et al. (Aug. 2018). “Detection of a Glitch in the Pulsar J1709-4429”. In: *Research Notes of the American Astronomical Society* 2.3, 139, p. 139. DOI: [10.3847/2515-5172/aad7bc](https://doi.org/10.3847/2515-5172/aad7bc). arXiv: [1808.02580](https://arxiv.org/abs/1808.02580) [astro-ph.HE].
- Lower, Marcus E. et al. (Dec. 2019). “Detection of a Glitch in PSR J0908-4913 by UTMOST”. In: *Research Notes of the American Astronomical Society* 3.12, 192, p. 192. DOI: [10.3847/2515-5172/ab621d](https://doi.org/10.3847/2515-5172/ab621d). arXiv: [1912.10827](https://arxiv.org/abs/1912.10827) [astro-ph.HE].

- Lower, Marcus E. et al. (Mar. 2023b). “The 2022 High-energy Outburst and Radio Disappearing Act of the Magnetar 1E 1547.0-5408”. In: *The Astrophysical Journal* 945.2, 153, p. 153. DOI: [10.3847/1538-4357/acbc7c](https://doi.org/10.3847/1538-4357/acbc7c). arXiv: [2302.07397](https://arxiv.org/abs/2302.07397) [astro-ph.HE].
- Lück, H. and GEO600 Team (June 1997). “The GEO600 project”. In: *Classical and Quantum Gravity* 14.6, pp. 1471–1476. DOI: [10.1088/0264-9381/14/6/012](https://doi.org/10.1088/0264-9381/14/6/012).
- Lyne, A. G. (July 1984). “Orbital inclination and mass of the binary pulsar PSR0655 + 64”. In: *Nature* 310.5975, pp. 300–302. DOI: [10.1038/310300a0](https://doi.org/10.1038/310300a0).
- (Apr. 1987). “A massive glitch in an old pulsar”. In: *Nature* 326.6113, pp. 569–571. DOI: [10.1038/326569a0](https://doi.org/10.1038/326569a0).
- Lyne, A. G., J. D. Biggs, P. A. Harrison, and M. Bailes (Jan. 1993). “A long-period globular-cluster pulsar in an eclipsing binary system”. In: *Nature* 361.6407, pp. 47–49. DOI: [10.1038/361047a0](https://doi.org/10.1038/361047a0).
- Lyne, A. G., C. A. Jordan, F. Graham-Smith, C. M. Espinoza, B. W. Stappers, and P. Weltevrede (Jan. 2015). “45 years of rotation of the Crab pulsar”. In: *Monthly Notices of the Royal Astronomical Society* 446.1, pp. 857–864. DOI: [10.1093/mnras/stu2118](https://doi.org/10.1093/mnras/stu2118). arXiv: [1410.0886](https://arxiv.org/abs/1410.0886) [astro-ph.HE].
- Lyne, A. G., R. S. Pritchard, F. Graham-Smith, and F. Camilo (June 1996). “Very low braking index for the Vela pulsar”. In: *Nature* 381.6582, pp. 497–498. DOI: [10.1038/381497a0](https://doi.org/10.1038/381497a0).
- Lyne, A. G., S. L. Shemar, and F. Graham Smith (July 2000). “Statistical studies of pulsar glitches”. In: *Monthly Notices of the Royal Astronomical Society* 315.3, pp. 534–542. DOI: [10.1046/j.1365-8711.2000.03415.x](https://doi.org/10.1046/j.1365-8711.2000.03415.x).
- Lyne, Andrew and Francis Graham-Smith (2012). *Pulsar Astronomy*.
- Lyne, Andrew, George Hobbs, Michael Kramer, Ingrid Stairs, and Ben Stappers (July 2010). “Switched Magnetospheric Regulation of Pulsar Spin-Down”. In: *Science* 329.5990, p. 408. DOI: [10.1126/science.1186683](https://doi.org/10.1126/science.1186683). arXiv: [1006.5184](https://arxiv.org/abs/1006.5184) [astro-ph.GA].
- Maggiore, Michele (Oct. 2007). *Gravitational Waves: Volume 1: Theory and Experiments*. Oxford University Press. ISBN: 9780198570745. DOI: [10.1093/acprof:oso/9780198570745.001.0001](https://doi.org/10.1093/acprof:oso/9780198570745.001.0001). URL: <https://doi.org/10.1093/acprof:oso/9780198570745.001.0001>.
- Maggiore, Michele et al. (Mar. 2020). “Science case for the Einstein telescope”. In: *Journal of Cosmology and Astroparticle Physics* 2020.3, 050, p. 050. DOI: [10.1088/1475-7516/2020/03/050](https://doi.org/10.1088/1475-7516/2020/03/050). arXiv: [1912.02622](https://arxiv.org/abs/1912.02622) [astro-ph.CO].
- Manchester, R. N. (Sept. 1971). “Observations of Pulsar Polarization at 410 and 1665 MHz”. In: *The Astrophysical Journal Supplement Series* 23, p. 283. DOI: [10.1086/190240](https://doi.org/10.1086/190240).
- Manchester, R. N., G. B. Hobbs, A. Teoh, and M. Hobbs (Apr. 2005). “The Australia Telescope National Facility Pulsar Catalogue”. In: *The Astronomical Journal* 129.4, pp. 1993–2006. DOI: [10.1086/428488](https://doi.org/10.1086/428488). arXiv: [astro-ph/0412641](https://arxiv.org/abs/astro-ph/0412641) [astro-ph].
- Manchester, R. N. and J. H. Taylor (July 1974). “Period Irregularities in Pulsars”. In: *The Astrophysical Journal* 191, p. L63. DOI: [10.1086/181549](https://doi.org/10.1086/181549).
- Mandlik, A. et al. (Aug. 2024). “The UTMOST-NS: a fully digital, wide-field transient search facility operating at a centre frequency of 831 MHz”. In: *Monthly Notices of the Royal Astronomical Society* 532.2, pp. 2644–2656. DOI: [10.1093/mnras/stae1309](https://doi.org/10.1093/mnras/stae1309).
- Mandlik, Ayushi (2024). “Unearthing Fast Radio Bursts using Machine Learning”. PhD thesis. Swinburne University of Technology.

- Maron, O., M. Serylak, J. Kijak, K. Krzeszowski, D. Mitra, and A. Jessner (July 2013). “Pulse-to-pulse flux density modulation from pulsars at 8.35 GHz”. In: *Astronomy and Astrophysics* 555, A28, A28. DOI: [10.1051/0004-6361/201220698](https://doi.org/10.1051/0004-6361/201220698). arXiv: [1304.6807](https://arxiv.org/abs/1304.6807) [astro-ph.SR].
- Marshall, F. E., E. V. Gotthelf, J. Middleditch, Q. D. Wang, and W. Zhang (Mar. 2004). “The Big Glitch: The Rotation History of PSR J0537-6910”. In: *The Astrophysical Journal* 603.2, pp. 682–689. DOI: [10.1086/381567](https://doi.org/10.1086/381567).
- Mastrano, A. and A. Melatos (Aug. 2005). “Kelvin-Helmholtz instability and circulation transfer at an isotropic-anisotropic superfluid interface in a neutron star”. In: *Monthly Notices of the Royal Astronomical Society* 361.3, pp. 927–941. DOI: [10.1111/j.1365-2966.2005.09219.x](https://doi.org/10.1111/j.1365-2966.2005.09219.x). arXiv: [astro-ph/0505529](https://arxiv.org/abs/astro-ph/0505529) [astro-ph].
- Matsakis, D. N., J. H. Taylor, and T. Marshall Eubanks (Oct. 1997). “A statistic for describing pulsar and clock stabilities.” In: *Astronomy and Astrophysics* 326, pp. 924–928.
- McCulloch, P. M., P. A. Hamilton, D. McConnell, and E. A. King (Aug. 1990). “The Vela glitch of Christmas 1988”. In: *Nature* 346.6287, pp. 822–824. DOI: [10.1038/346822a0](https://doi.org/10.1038/346822a0).
- McCulloch, P. M., A. R. Klekociuk, P. A. Hamilton, and G. W. R. Royle (Jan. 1987). “Daily observations of three period jumps of the VELA pulsar.” In: *Australian Journal of Physics* 40, pp. 725–730. DOI: [10.1071/PH870725](https://doi.org/10.1071/PH870725).
- McKee, J. W. et al. (Sept. 2016). “A glitch in the millisecond pulsar J0613-0200”. In: *Monthly Notices of the Royal Astronomical Society* 461.3, pp. 2809–2817. DOI: [10.1093/mnras/stw1442](https://doi.org/10.1093/mnras/stw1442). arXiv: [1606.04098](https://arxiv.org/abs/1606.04098) [astro-ph.HE].
- McKenna, J. and A. G. Lyne (Jan. 1990). “PSR1737-30 and period discontinuities in young pulsars”. In: *Nature* 343.6256, pp. 349–350. DOI: [10.1038/343349a0](https://doi.org/10.1038/343349a0).
- McKinnon, Mark M. and Daniel R. Stinebring (Jan. 2000). “The Mode-separated Pulse Profiles of Pulsar Radio Emission”. In: *The Astrophysical Journal* 529.1, pp. 435–446. DOI: [10.1086/308264](https://doi.org/10.1086/308264).
- Melatos, A. (July 1997). “Spin-down of an oblique rotator with a current-starved outer magnetosphere”. In: *Monthly Notices of the Royal Astronomical Society* 288.4, pp. 1049–1059. DOI: [10.1093/mnras/288.4.1049](https://doi.org/10.1093/mnras/288.4.1049).
- Melatos, A., P. Clearwater, S. Suvorova, L. Sun, W. Moran, and R. J. Evans (Aug. 2021). “Hidden Markov model tracking of continuous gravitational waves from a binary neutron star with wandering spin. III. Rotational phase tracking”. In: *Physical Review D* 104.4, 042003, p. 042003. DOI: [10.1103/PhysRevD.104.042003](https://doi.org/10.1103/PhysRevD.104.042003). arXiv: [2107.12822](https://arxiv.org/abs/2107.12822) [gr-qc].
- Melatos, A., L. M. Dunn, S. Suvorova, W. Moran, and R. J. Evans (June 2020). “Pulsar Glitch Detection with a Hidden Markov Model”. In: *The Astrophysical Journal* 896.1, 78, p. 78. DOI: [10.3847/1538-4357/ab9178](https://doi.org/10.3847/1538-4357/ab9178). arXiv: [2005.09388](https://arxiv.org/abs/2005.09388) [astro-ph.HE].
- Melatos, A., G. Howitt, and W. Fulgenzi (Aug. 2018). “Size-waiting-time Correlations in Pulsar Glitches”. In: *The Astrophysical Journal* 863.2, 196, p. 196. DOI: [10.3847/1538-4357/aad228](https://doi.org/10.3847/1538-4357/aad228). arXiv: [1809.03064](https://arxiv.org/abs/1809.03064) [astro-ph.HE].
- Melatos, A., N. J. O’Neill, P. M. Meyers, and J. O’Leary (Feb. 2023). “Tracking Hidden Magnetospheric Fluctuations in Accretion-powered Pulsars With a Kalman Filter”. In: *The Astrophysical Journal* 944.1, 64, p. 64. DOI: [10.3847/1538-4357/acab5a](https://doi.org/10.3847/1538-4357/acab5a). arXiv: [2212.08185](https://arxiv.org/abs/2212.08185) [astro-ph.HE].

- Melatos, A. and D. J. B. Payne (Apr. 2005). “Gravitational Radiation from an Accreting Millisecond Pulsar with a Magnetically Confined Mountain”. In: *The Astrophysical Journal* 623.2, pp. 1044–1050. DOI: [10.1086/428600](https://doi.org/10.1086/428600). arXiv: [astro-ph/0503287](https://arxiv.org/abs/astro-ph/0503287) [astro-ph].
- Melatos, A., C. Peralta, and J. S. B. Wyithe (Jan. 2008). “Avalanche Dynamics of Radio Pulsar Glitches”. In: *The Astrophysical Journal* 672.2, pp. 1103–1118. DOI: [10.1086/523349](https://doi.org/10.1086/523349). arXiv: [0710.1021](https://arxiv.org/abs/0710.1021) [astro-ph].
- Melatos, A. and E. S. Phinney (Jan. 2001). “Hydromagnetic Structure of a Neutron Star Accreting at Its Polar Caps”. In: *Publications of the Astronomical Society of Australia* 18.4, pp. 421–430. DOI: [10.1071/AS01056](https://doi.org/10.1071/AS01056).
- Melatos, Andrew and Bennett Link (Jan. 2014). “Pulsar timing noise from superfluid turbulence”. In: *Monthly Notices of the Royal Astronomical Society* 437.1, pp. 21–31. DOI: [10.1093/mnras/stt1828](https://doi.org/10.1093/mnras/stt1828). arXiv: [1310.3108](https://arxiv.org/abs/1310.3108) [astro-ph.HE].
- Messenger, C. et al. (July 2015). “Gravitational waves from Scorpius X-1: A comparison of search methods and prospects for detection with advanced detectors”. In: *Physical Review D* 92.2, 023006, p. 023006. DOI: [10.1103/PhysRevD.92.023006](https://doi.org/10.1103/PhysRevD.92.023006). arXiv: [1504.05889](https://arxiv.org/abs/1504.05889) [gr-qc].
- Meyers, Patrick M., Andrew Melatos, and Nicholas J. O’Neill (Apr. 2021a). “Parameter estimation of a two-component neutron star model with spin wandering”. In: *Monthly Notices of the Royal Astronomical Society* 502.3, pp. 3113–3127. DOI: [10.1093/mnras/stab262](https://doi.org/10.1093/mnras/stab262). arXiv: [2101.12421](https://arxiv.org/abs/2101.12421) [astro-ph.HE].
- Meyers, Patrick M., Nicholas J. O’Neill, Andrew Melatos, and Robin J. Evans (Sept. 2021b). “Rapid parameter estimation of a two-component neutron star model with spin wandering using a Kalman filter”. In: *Monthly Notices of the Royal Astronomical Society* 506.3, pp. 3349–3363. DOI: [10.1093/mnras/stab1952](https://doi.org/10.1093/mnras/stab1952). arXiv: [2107.03047](https://arxiv.org/abs/2107.03047) [astro-ph.HE].
- Middleditch, J., F. E. Marshall, Q. D. Wang, E. V. Gotthelf, and W. Zhang (Dec. 2006). “Predicting the Starquakes in PSR J0537-6910”. In: *The Astrophysical Journal* 652.2, pp. 1531–1546. DOI: [10.1086/508736](https://doi.org/10.1086/508736). arXiv: [astro-ph/0605007](https://arxiv.org/abs/astro-ph/0605007) [astro-ph].
- Middleton, Hannah, Patrick Clearwater, Andrew Melatos, and Liam Dunn (July 2020). “Search for gravitational waves from five low mass x-ray binaries in the second Advanced LIGO observing run with an improved hidden Markov model”. In: *Physical Review D* 102.2, 023006, p. 023006. DOI: [10.1103/PhysRevD.102.023006](https://doi.org/10.1103/PhysRevD.102.023006). arXiv: [2006.06907](https://arxiv.org/abs/2006.06907) [astro-ph.HE].
- Millhouse, M., A. Melatos, G. Howitt, J. B. Carlin, L. Dunn, and G. Ashton (Apr. 2022). “An updated glitch rate law inferred from radio pulsars”. In: *Monthly Notices of the Royal Astronomical Society* 511.3, pp. 3304–3319. DOI: [10.1093/mnras/stac194](https://doi.org/10.1093/mnras/stac194). arXiv: [2202.01930](https://arxiv.org/abs/2202.01930) [astro-ph.HE].
- Millhouse, Margaret, Lucy Strang, and Andrew Melatos (Oct. 2020). “Search for gravitational waves from 12 young supernova remnants with a hidden Markov model in Advanced LIGO’s second observing run”. In: *Physical Review D* 102.8, 083025, p. 083025. DOI: [10.1103/PhysRevD.102.083025](https://doi.org/10.1103/PhysRevD.102.083025). arXiv: [2003.08588](https://arxiv.org/abs/2003.08588) [gr-qc].
- Mongiòvi, Maria Stella, Francesco G. Russo, and Michele Sciacca (Aug. 2017). “A mathematical description of glitches in neutron stars”. In: *Monthly Notices of the Royal Astronomical Society* 469.2, pp. 2141–2150. DOI: [10.1093/mnras/stx827](https://doi.org/10.1093/mnras/stx827).

- Montoli, A., M. Antonelli, and P. M. Pizzochero (Mar. 2020). “The role of mass, equation of state, and superfluid reservoir in large pulsar glitches”. In: *Monthly Notices of the Royal Astronomical Society* 492.4, pp. 4837–4846. DOI: [10.1093/mnras/staa14910.1038/ncomms1195](https://doi.org/10.1093/mnras/staa14910.1038/ncomms1195). arXiv: [1809.07834](https://arxiv.org/abs/1809.07834) [astro-ph.HE].
- Montoli, Alessandro, Marco Antonelli, Brynmor Haskell, and Pierre Pizzochero (Jan. 2021). “Statistical Estimates of the Pulsar Glitch Activity”. In: *Universe* 7.1, 8, p. 8. DOI: [10.3390/universe7010008](https://doi.org/10.3390/universe7010008). arXiv: [2012.01539](https://arxiv.org/abs/2012.01539) [astro-ph.HE].
- Morales, J. A. and C. J. Horowitz (Dec. 2022). “Neutron star crust can support a large ellipticity”. In: *Monthly Notices of the Royal Astronomical Society* 517.4, pp. 5610–5616. DOI: [10.1093/mnras/stac3058](https://doi.org/10.1093/mnras/stac3058). arXiv: [2209.03222](https://arxiv.org/abs/2209.03222) [gr-qc].
- Mukherjee, Arunava, Chris Messenger, and Keith Riles (Feb. 2018). “Accretion-induced spin-wandering effects on the neutron star in Scorpius X-1: Implications for continuous gravitational wave searches”. In: *Physical Review D* 97.4, 043016, p. 043016. DOI: [10.1103/PhysRevD.97.043016](https://doi.org/10.1103/PhysRevD.97.043016). arXiv: [1710.06185](https://arxiv.org/abs/1710.06185) [gr-qc].
- Mukherjee, Dipanjan (Sept. 2017). “Revisiting Field Burial by Accretion onto Neutron Stars”. In: *Journal of Astrophysics and Astronomy* 38.3, 48, p. 48. DOI: [10.1007/s12036-017-9465-6](https://doi.org/10.1007/s12036-017-9465-6). arXiv: [1709.07332](https://arxiv.org/abs/1709.07332) [astro-ph.HE].
- Mukherjee, Dipanjan, Dipankar Bhattacharya, and Andrea Mignone (Apr. 2013). “MHD instabilities in accretion mounds - I. 2D axisymmetric simulations”. In: *Monthly Notices of the Royal Astronomical Society* 430.3, pp. 1976–1987. DOI: [10.1093/mnras/stt020](https://doi.org/10.1093/mnras/stt020). arXiv: [1212.3897](https://arxiv.org/abs/1212.3897) [astro-ph.HE].
- Nelson, Jerry, Richard Hills, David Cudaback, and Joseph Wampler (Sept. 1970). “Optical Timing of the Pulsar NP 0532 IN the Crab Nebula”. In: *The Astrophysical Journal* 161, p. L235. DOI: [10.1086/180609](https://doi.org/10.1086/180609).
- Newville, Matthew, Till Stensitzki, Daniel B. Allen, and Antonino Ingargiola (Sept. 2014). *LMFIT: Non-Linear Least-Square Minimization and Curve-Fitting for Python*. Version 0.8.0. DOI: [10.5281/zenodo.11813](https://doi.org/10.5281/zenodo.11813).
- Newville, Matthew, Till Stensitzki, Daniel B. Allen, Michal Rawlik, Antonino Ingargiola, and Andrew Nelson (June 2016). *Lmfit: Non-Linear Least-Square Minimization and Curve-Fitting for Python*. Astrophysics Source Code Library, record ascl:1606.014.
- Ng, Cherry (Aug. 2018). “Pulsar science with the CHIME telescope”. In: *Pulsar Astrophysics the Next Fifty Years*. Ed. by P. Weltevrede, B. B. P. Perera, L. L. Preston, and S. Sanidas. Vol. 337. IAU Symposium, pp. 179–182. DOI: [10.1017/S1743921317010638](https://doi.org/10.1017/S1743921317010638). arXiv: [1711.02104](https://arxiv.org/abs/1711.02104) [astro-ph.IM].
- Nițu, Iuliana C., Michael J. Keith, Ben W. Stappers, Andrew G. Lyne, and Mitchell B. Mickaliger (May 2022). “A search for planetary companions around 800 pulsars from the Jodrell Bank pulsar timing programme”. In: *Monthly Notices of the Royal Astronomical Society* 512.2, pp. 2446–2459. DOI: [10.1093/mnras/stac593](https://doi.org/10.1093/mnras/stac593). arXiv: [2203.01136](https://arxiv.org/abs/2203.01136) [astro-ph.EP].
- Nieder, L. et al. (Sept. 2019). “Detection and Timing of Gamma-Ray Pulsations from the 707 Hz Pulsar J0952-0607”. In: *The Astrophysical Journal* 883.1, 42, p. 42. DOI: [10.3847/1538-4357/ab357e](https://doi.org/10.3847/1538-4357/ab357e). arXiv: [1905.11352](https://arxiv.org/abs/1905.11352) [astro-ph.HE].
- Olausen, S. A. and V. M. Kaspi (May 2014). “The McGill Magnetar Catalog”. In: *The Astrophysical Journal Supplement Series* 212.1, 6, p. 6. DOI: [10.1088/0067-0049/212/1/6](https://doi.org/10.1088/0067-0049/212/1/6). arXiv: [1309.4167](https://arxiv.org/abs/1309.4167) [astro-ph.HE].

- O’Leary, Joseph, Andrew Melatos, Tom Kimpson, Nicholas J. O’Neill, Patrick M. Meyers, Dimitris M. Christodoulou, Sayantan Bhattacharya, and Silas G. T. Laycock (Aug. 2024a). “Measuring the Magnetic Dipole Moment and Magnetospheric Fluctuations of Accretion-powered Pulsars in the Small Magellanic Cloud with an Unscented Kalman Filter”. In: *The Astrophysical Journal* 971.2, 126, p. 126. DOI: [10.3847/1538-4357/ad53c2](https://doi.org/10.3847/1538-4357/ad53c2). arXiv: [2406.01827](https://arxiv.org/abs/2406.01827) [astro-ph.HE].
- O’Leary, Joseph, Andrew Melatos, Nicholas J. O’Neill, Patrick M. Meyers, Dimitris M. Christodoulou, Sayantan Bhattacharya, and Silas G. T. Laycock (Apr. 2024b). “Measuring the Magnetic Dipole Moment and Magnetospheric Fluctuations of SXP 18.3 with a Kalman Filter”. In: *The Astrophysical Journal* 965.2, 102, p. 102. DOI: [10.3847/1538-4357/ad2adc](https://doi.org/10.3847/1538-4357/ad2adc). arXiv: [2402.11991](https://arxiv.org/abs/2402.11991) [astro-ph.HE].
- O’Leary, Joseph et al. (in prep.). “Discovering binary pulsars with a hidden Markov model”.
- Olney, Steve (July 2021). “Glitch event in the Vela pulsar (PSR J0835-4510) observed at HawkRAO”. in: *The Astronomer’s Telegram* 14808, p. 1.
- O’Neill, Nicholas J., Patrick M. Meyers, and Andrew Melatos (June 2024). “Analysing radio pulsar timing noise with a Kalman filter: a demonstration involving PSR J1359-6038”. In: *Monthly Notices of the Royal Astronomical Society* 530.4, pp. 4648–4664. DOI: [10.1093/mnras/stae770](https://doi.org/10.1093/mnras/stae770). arXiv: [2403.16467](https://arxiv.org/abs/2403.16467) [astro-ph.HE].
- Onuchukwu, C. C. and E. Legahara (May 2024). “Comparative analysis of the parameters of pulsars with braking indices $n > 0$ and $n < 0$ ”. In: *Astrophysics and Space Science* 369.5, 51, p. 51. DOI: [10.1007/s10509-024-04317-3](https://doi.org/10.1007/s10509-024-04317-3).
- Osborne, E. L. and D. I. Jones (May 2020). “Gravitational waves from magnetically induced thermal neutron star mountains”. In: *Monthly Notices of the Royal Astronomical Society* 494.2, pp. 2839–2850. DOI: [10.1093/mnras/staa858](https://doi.org/10.1093/mnras/staa858). arXiv: [1910.04453](https://arxiv.org/abs/1910.04453) [astro-ph.HE].
- Osłowski, S., W. van Straten, M. Bailes, A. Jameson, and G. Hobbs (July 2014). “Timing, polarimetry and physics of the bright, nearby millisecond pulsar PSR J0437-4715 - a single-pulse perspective”. In: *Monthly Notices of the Royal Astronomical Society* 441.4, pp. 3148–3160. DOI: [10.1093/mnras/stu804](https://doi.org/10.1093/mnras/stu804). arXiv: [1405.2638](https://arxiv.org/abs/1405.2638) [astro-ph.SR].
- Oswald, Lucy, Aris Karastergiou, and Simon Johnston (Aug. 2020). “Pulsar polarimetry with the Parkes ultra-wideband receiver”. In: *Monthly Notices of the Royal Astronomical Society* 496.2, pp. 1418–1429. DOI: [10.1093/mnras/staa1597](https://doi.org/10.1093/mnras/staa1597). arXiv: [2006.04726](https://arxiv.org/abs/2006.04726) [astro-ph.HE].
- Owen, Benjamin J. (Nov. 2010). “How to adapt broad-band gravitational-wave searches for r-modes”. In: *Physical Review D* 82.10, 104002, p. 104002. DOI: [10.1103/PhysRevD.82.104002](https://doi.org/10.1103/PhysRevD.82.104002). arXiv: [1006.1994](https://arxiv.org/abs/1006.1994) [gr-qc].
- Owen, Benjamin J., Lee Lindblom, Curt Cutler, Bernard F. Schutz, Alberto Vecchio, and Nils Andersson (Oct. 1998). “Gravitational waves from hot young rapidly rotating neutron stars”. In: *Physical Review D* 58.8, 084020, p. 084020. DOI: [10.1103/PhysRevD.58.084020](https://doi.org/10.1103/PhysRevD.58.084020). arXiv: [gr-qc/9804044](https://arxiv.org/abs/gr-qc/9804044) [gr-qc].
- Pacini, F. (Nov. 1967). “Energy Emission from a Neutron Star”. In: *Nature* 216.5115, pp. 567–568. DOI: [10.1038/216567a0](https://doi.org/10.1038/216567a0).
- Palfreyman, Jim (Dec. 2016). “Glitch observed in the Vela Pulsar (PSR J0835-4510)”. In: *The Astronomer’s Telegram* 9847, p. 1.

- Palfreyman, Jim, John M. Dickey, Aidan Hotan, Simon Ellingsen, and Willem van Straten (Apr. 2018). “Alteration of the magnetosphere of the Vela pulsar during a glitch”. In: *Nature* 556.7700, pp. 219–222. DOI: [10.1038/s41586-018-0001-x](https://doi.org/10.1038/s41586-018-0001-x).
- Palomba, C. (Feb. 2001). “Gravitational radiation from young magnetars: Preliminary results”. In: *Astronomy and Astrophysics* 367, pp. 525–531. DOI: [10.1051/0004-6361:20000452](https://doi.org/10.1051/0004-6361:20000452).
- (May 2005). “Simulation of a population of isolated neutron stars evolving through the emission of gravitational waves”. In: *Monthly Notices of the Royal Astronomical Society* 359.3, pp. 1150–1164. DOI: [10.1111/j.1365-2966.2005.08975.x](https://doi.org/10.1111/j.1365-2966.2005.08975.x). arXiv: [astro-ph/0503046](https://arxiv.org/abs/astro-ph/0503046) [[astro-ph](#)].
- Papitto, A. et al. (Sept. 2013). “Swings between rotation and accretion power in a binary millisecond pulsar”. In: *Nature* 501.7468, pp. 517–520. DOI: [10.1038/nature12470](https://doi.org/10.1038/nature12470). arXiv: [1305.3884](https://arxiv.org/abs/1305.3884) [[astro-ph.HE](#)].
- Parthasarathy, A. et al. (Nov. 2019). “Timing of young radio pulsars - I. Timing noise, periodic modulation, and proper motion”. In: *Monthly Notices of the Royal Astronomical Society* 489.3, pp. 3810–3826. DOI: [10.1093/mnras/stz2383](https://doi.org/10.1093/mnras/stz2383). arXiv: [1908.11709](https://arxiv.org/abs/1908.11709) [[astro-ph.HE](#)].
- Parthasarathy, A. et al. (May 2020). “Timing of young radio pulsars - II. Braking indices and their interpretation”. In: *Monthly Notices of the Royal Astronomical Society* 494.2, pp. 2012–2026. DOI: [10.1093/mnras/staa882](https://doi.org/10.1093/mnras/staa882). arXiv: [2003.13303](https://arxiv.org/abs/2003.13303) [[astro-ph.HE](#)].
- Parthasarathy, A. et al. (Mar. 2021). “Measurements of pulse jitter and single-pulse variability in millisecond pulsars using MeerKAT”. In: *Monthly Notices of the Royal Astronomical Society* 502.1, pp. 407–422. DOI: [10.1093/mnras/stab037](https://doi.org/10.1093/mnras/stab037). arXiv: [2101.08531](https://arxiv.org/abs/2101.08531) [[astro-ph.HE](#)].
- Patruno, Alessandro and Anna L. Watts (Jan. 2021). “Accreting Millisecond X-ray Pulsars”. In: *Timing Neutron Stars: Pulsations, Oscillations and Explosions*. Ed. by Tomaso M. Belloni, Mariano Méndez, and Chengmin Zhang. Vol. 461. Astrophysics and Space Science Library, pp. 143–208. DOI: [10.1007/978-3-662-62110-3_4](https://doi.org/10.1007/978-3-662-62110-3_4). arXiv: [1206.2727](https://arxiv.org/abs/1206.2727) [[astro-ph.HE](#)].
- Payne, D. J. B. and A. Melatos (June 2004). “Burial of the polar magnetic field of an accreting neutron star - I. Self-consistent analytic and numerical equilibria”. In: *Monthly Notices of the Royal Astronomical Society* 351.2, pp. 569–584. DOI: [10.1111/j.1365-2966.2004.07798.x](https://doi.org/10.1111/j.1365-2966.2004.07798.x). arXiv: [astro-ph/0403173](https://arxiv.org/abs/astro-ph/0403173) [[astro-ph](#)].
- (Apr. 2006). “Frequency Spectrum of Gravitational Radiation from Global Hydro-magnetic Oscillations of a Magnetically Confined Mountain on an Accreting Neutron Star”. In: *The Astrophysical Journal* 641.1, pp. 471–478. DOI: [10.1086/498855](https://doi.org/10.1086/498855). arXiv: [astro-ph/0510053](https://arxiv.org/abs/astro-ph/0510053) [[astro-ph](#)].
- Peralta, C., A. Melatos, M. Giacobello, and A. Ooi (Nov. 2006). “Transitions between Turbulent and Laminar Superfluid Vorticity States in the Outer Core of a Neutron Star”. In: *The Astrophysical Journal* 651.2, pp. 1079–1091. DOI: [10.1086/507576](https://doi.org/10.1086/507576). arXiv: [astro-ph/0607161](https://arxiv.org/abs/astro-ph/0607161) [[astro-ph](#)].
- Peralta, Carlos Andrés (2006). “Superfluid spherical Couette flow and rotational irregularities in pulsars”. PhD thesis. University of Melbourne.
- Philippov, A. and M. Kramer (Aug. 2022). “Pulsar Magnetospheres and Their Radiation”. In: *Annual Review of Astronomy and Astrophysics* 60, pp. 495–558. DOI: [10.1146/annurev-astro-052920-112338](https://doi.org/10.1146/annurev-astro-052920-112338).

- Phinney, E. S. (Oct. 1992). “Pulsars as Probes of Newtonian Dynamical Systems”. In: *Philosophical Transactions of the Royal Society of London Series A* 341.1660, pp. 39–75. DOI: [10.1098/rsta.1992.0084](https://doi.org/10.1098/rsta.1992.0084).
- Piccinni, Ornella J. et al. (Apr. 2020). “Directed search for continuous gravitational-wave signals from the Galactic Center in the Advanced LIGO second observing run”. In: *Physical Review D* 101.8, 082004, p. 082004. DOI: [10.1103/PhysRevD.101.082004](https://doi.org/10.1103/PhysRevD.101.082004). arXiv: [1910.05097](https://arxiv.org/abs/1910.05097) [gr-qc].
- Piccinni, Ornella Juliana (May 2022). “Status and Perspectives of Continuous Gravitational Wave Searches”. In: *Galaxies* 10.3, 72, p. 72. DOI: [10.3390/galaxies10030072](https://doi.org/10.3390/galaxies10030072). arXiv: [2202.01088](https://arxiv.org/abs/2202.01088) [gr-qc].
- Pines, David and Jacob Shaham (Jan. 1972). “Microquakes and Macroquakes in Neutron Stars”. In: *Nature Physical Science* 235.55, pp. 43–49. DOI: [10.1038/physci235043a0](https://doi.org/10.1038/physci235043a0).
- Pitkin, M., C. Gill, D. I. Jones, G. Woan, and G. S. Davies (Nov. 2015). “First results and future prospects for dual-harmonic searches for gravitational waves from spinning neutron stars”. In: *Monthly Notices of the Royal Astronomical Society* 453.4, pp. 4399–4420. DOI: [10.1093/mnras/stv1931](https://doi.org/10.1093/mnras/stv1931). arXiv: [1508.00416](https://arxiv.org/abs/1508.00416) [astro-ph.HE].
- Pitkin, Matthew (Feb. 2018). “psrqpy: a python interface for querying the ATNF pulsar catalogue”. In: *The Journal of Open Source Software* 3.22, p. 538. DOI: [10.21105/joss.00538](https://doi.org/10.21105/joss.00538). arXiv: [1806.07809](https://arxiv.org/abs/1806.07809) [astro-ph.IM].
- Pitkin, Matthew, Maximiliano Isi, John Veitch, and Graham Woan (May 2017). “A nested sampling code for targeted searches for continuous gravitational waves from pulsars”. In: *arXiv e-prints*, arXiv:1705.08978, arXiv:1705.08978. DOI: [10.48550/arXiv.1705.08978](https://doi.org/10.48550/arXiv.1705.08978). arXiv: [1705.08978](https://arxiv.org/abs/1705.08978) [gr-qc].
- Pletsch, Holger J. (Aug. 2010). “Parameter-space metric of semicoherent searches for continuous gravitational waves”. In: *Physical Review D* 82.4, 042002, p. 042002. DOI: [10.1103/PhysRevD.82.042002](https://doi.org/10.1103/PhysRevD.82.042002). arXiv: [1005.0395](https://arxiv.org/abs/1005.0395) [gr-qc].
- Pooley, David et al. (July 2003). “Dynamical Formation of Close Binary Systems in Globular Clusters”. In: *The Astrophysical Journal* 591.2, pp. L131–L134. DOI: [10.1086/377074](https://doi.org/10.1086/377074). arXiv: [astro-ph/0305003](https://arxiv.org/abs/astro-ph/0305003) [astro-ph].
- Preskill, John, Mark B. Wise, and Frank Wilczek (Jan. 1983). “Cosmology of the invisible axion”. In: *Physics Letters B* 120.1-3, pp. 127–132. DOI: [10.1016/0370-2693\(83\)90637-8](https://doi.org/10.1016/0370-2693(83)90637-8).
- Prix, R. and M. Shaltev (Apr. 2012). “Search for continuous gravitational waves: Optimal StackSlide method at fixed computing cost”. In: *Physical Review D* 85.8, 084010, p. 084010. DOI: [10.1103/PhysRevD.85.084010](https://doi.org/10.1103/PhysRevD.85.084010). arXiv: [1201.4321](https://arxiv.org/abs/1201.4321) [gr-qc].
- Prix, Reinhard (Jan. 2007a). “Search for continuous gravitational waves: Metric of the multidetector F-statistic”. In: *Physical Review D* 75.2, 023004, p. 023004. DOI: [10.1103/PhysRevD.75.023004](https://doi.org/10.1103/PhysRevD.75.023004). arXiv: [gr-qc/0606088](https://arxiv.org/abs/gr-qc/0606088) [gr-qc].
- (Oct. 2007b). “Template-based searches for gravitational waves: efficient lattice covering of flat parameter spaces”. In: *Classical and Quantum Gravity* 24.19, S481–S490. DOI: [10.1088/0264-9381/24/19/S11](https://doi.org/10.1088/0264-9381/24/19/S11). arXiv: [0707.0428](https://arxiv.org/abs/0707.0428) [gr-qc].
- Prix, Reinhard and Yousuke Itoh (Sept. 2005). “Global parameter-space correlations of coherent searches for continuous gravitational waves”. In: *Classical and Quantum Gravity* 22.18, S1003–S1012. DOI: [10.1088/0264-9381/22/18/S14](https://doi.org/10.1088/0264-9381/22/18/S14). arXiv: [gr-qc/0504006](https://arxiv.org/abs/gr-qc/0504006) [gr-qc].

- Prix, Reinhard and Badri Krishnan (Oct. 2009). “Targeted search for continuous gravitational waves: Bayesian versus maximum-likelihood statistics”. In: *Classical and Quantum Gravity* 26.20, 204013, p. 204013. DOI: [10.1088/0264-9381/26/20/204013](https://doi.org/10.1088/0264-9381/26/20/204013). arXiv: [0907.2569](https://arxiv.org/abs/0907.2569) [gr-qc].
- Priymak, M., A. Melatos, and D. J. B. Payne (Nov. 2011). “Quadrupole moment of a magnetically confined mountain on an accreting neutron star: effect of the equation of state”. In: *Monthly Notices of the Royal Astronomical Society* 417.4, pp. 2696–2713. DOI: [10.1111/j.1365-2966.2011.19431.x](https://doi.org/10.1111/j.1365-2966.2011.19431.x). arXiv: [1109.1040](https://arxiv.org/abs/1109.1040) [astro-ph.HE].
- Rabiner, L.R. (1989). “A tutorial on hidden Markov models and selected applications in speech recognition”. In: *Proceedings of the IEEE* 77.2, pp. 257–286. DOI: [10.1109/5.18626](https://doi.org/10.1109/5.18626).
- Radhakrishnan, V. and R. N. Manchester (Apr. 1969). “Detection of a Change of State in the Pulsar PSR 0833-45”. In: *Nature* 222.5190, pp. 228–229. DOI: [10.1038/222228a0](https://doi.org/10.1038/222228a0).
- Ransom, Scott (July 2011). *PRESTO: Pulsar Exploration and Search TOolkit*. Astrophysics Source Code Library, record ascl:1107.017.
- Ransom, Scott M. (May 2008). “Pulsars in Globular Clusters”. In: *Dynamical Evolution of Dense Stellar Systems*. Ed. by Enrico Vesperini, Mirek Giersz, and Alison Sills. Vol. 246. IAU Symposium, pp. 291–300. DOI: [10.1017/S1743921308015810](https://doi.org/10.1017/S1743921308015810).
- Ransom, Scott M., James M. Cordes, and Stephen S. Eikenberry (June 2003). “A New Search Technique for Short Orbital Period Binary Pulsars”. In: *The Astrophysical Journal* 589.2, pp. 911–920. DOI: [10.1086/374806](https://doi.org/10.1086/374806). arXiv: [astro-ph/0210010](https://arxiv.org/abs/astro-ph/0210010) [astro-ph].
- Ransom, Scott Mitchell (Jan. 2001). “New search techniques for binary pulsars”. PhD thesis. Harvard University, Massachusetts.
- Reardon, Daniel J. et al. (July 2023). “Search for an Isotropic Gravitational-wave Background with the Parkes Pulsar Timing Array”. In: *The Astrophysical Journal* 951.1, L6, p. L6. DOI: [10.3847/2041-8213/acdd02](https://doi.org/10.3847/2041-8213/acdd02). arXiv: [2306.16215](https://arxiv.org/abs/2306.16215) [astro-ph.HE].
- Reardon, Daniel John (Sept. 2018). “Precision Radio-Frequency Pulsar Timing & Interstellar Scintillometry”. PhD thesis. Monash University, School of Astrophysics.
- Reichley, P. E. and G. S. Downs (Apr. 1969). “Observed Decrease in the Periods of Pulsar PSR 0833-45”. In: *Nature* 222.5190, pp. 229–230. DOI: [10.1038/222229a0](https://doi.org/10.1038/222229a0).
- Reig, Pablo (Mar. 2011). “Be/X-ray binaries”. In: *Astrophysics and Space Science* 332.1, pp. 1–29. DOI: [10.1007/s10509-010-0575-8](https://doi.org/10.1007/s10509-010-0575-8). arXiv: [1101.5036](https://arxiv.org/abs/1101.5036) [astro-ph.HE].
- Rencoret, Javier A., Claudia Aguilera-Gómez, and Andreas Reisenegger (Oct. 2021). “Revisiting neutron starquakes caused by spin-down”. In: *Astronomy and Astrophysics* 654, A47, A47. DOI: [10.1051/0004-6361/202141499](https://doi.org/10.1051/0004-6361/202141499). arXiv: [2106.12604](https://arxiv.org/abs/2106.12604) [astro-ph.HE].
- Ridolfi, A. et al. (Aug. 2022). “TRAPUM discovery of 13 new pulsars in NGC 1851 using MeerKAT”. in: *Astronomy and Astrophysics* 664, A27, A27. DOI: [10.1051/0004-6361/202143006](https://doi.org/10.1051/0004-6361/202143006). arXiv: [2203.12302](https://arxiv.org/abs/2203.12302) [astro-ph.HE].
- Riggio, A., T. Di Salvo, L. Burderi, M. T. Menna, A. Papitto, R. Iaria, and G. Lavagetto (May 2008). “Spin-up and Phase Fluctuations in the Timing of the Accreting Millisecond Pulsar XTE J1807-294”. In: *The Astrophysical Journal* 678.2, pp. 1273–1278. DOI: [10.1086/533578](https://doi.org/10.1086/533578). arXiv: [0710.3450](https://arxiv.org/abs/0710.3450) [astro-ph].
- Riggio, A., A. Papitto, L. Burderi, T. di Salvo, M. Bachetti, R. Iaria, A. D’Ai, and M. T. Menna (Feb. 2011). “Timing of the accreting millisecond pulsar IGR J17511-3057”. In: *Astronomy and Astrophysics* 526, A95, A95. DOI: [10.1051/0004-6361/201014322](https://doi.org/10.1051/0004-6361/201014322). arXiv: [1010.6186](https://arxiv.org/abs/1010.6186) [astro-ph.HE].

- Riles, Keith (Dec. 2023). “Searches for continuous-wave gravitational radiation”. In: *Living Reviews in Relativity* 26.1, 3, p. 3. DOI: [10.1007/s41114-023-00044-3](https://doi.org/10.1007/s41114-023-00044-3). arXiv: [2206.06447](https://arxiv.org/abs/2206.06447) [astro-ph.HE].
- Rossetto, Pedro H. B., Jörg Frauendiener, Ryan Brunet, and Andrew Melatos (Dec. 2023). “Magnetically confined mountains on accreting neutron stars in general relativity”. In: *Monthly Notices of the Royal Astronomical Society* 526.2, pp. 2058–2066. DOI: [10.1093/mnras/stad2850](https://doi.org/10.1093/mnras/stad2850). arXiv: [2309.09519](https://arxiv.org/abs/2309.09519) [astro-ph.HE].
- Ruderman, M. (Aug. 1969). “Neutron Starquakes and Pulsar Periods”. In: *Nature* 223.5206, pp. 597–598. DOI: [10.1038/223597b0](https://doi.org/10.1038/223597b0).
- (Jan. 1976). “Crust-breaking by neutron superfluids and the Vela pulsar glitches.” In: *The Astrophysical Journal* 203, pp. 213–222. DOI: [10.1086/154069](https://doi.org/10.1086/154069).
- Ruderman, M. A. and P. G. Sutherland (Feb. 1975). “Theory of pulsars: polar gaps, sparks, and coherent microwave radiation.” In: *The Astrophysical Journal* 196, pp. 51–72. DOI: [10.1086/153393](https://doi.org/10.1086/153393).
- Ruderman, Malvin, Tianhua Zhu, and Kaiyou Chen (Jan. 1998). “Neutron Star Magnetic Field Evolution, Crust Movement, and Glitches”. In: *The Astrophysical Journal* 492.1, pp. 267–280. DOI: [10.1086/305026](https://doi.org/10.1086/305026). arXiv: [astro-ph/9709008](https://arxiv.org/abs/astro-ph/9709008) [astro-ph].
- Rüster, Stefan B., Matthias Hempel, and Jürgen Schaffner-Bielich (Mar. 2006). “Outer crust of nonaccreting cold neutron stars”. In: *Physical Review C* 73.3, 035804, p. 035804. DOI: [10.1103/PhysRevC.73.035804](https://doi.org/10.1103/PhysRevC.73.035804). arXiv: [astro-ph/0509325](https://arxiv.org/abs/astro-ph/0509325) [astro-ph].
- Santiago-Prieto, I., I. S. Heng, D. I. Jones, and J. Clark (June 2012). “Prospects for transient gravitational waves at r-mode frequencies associated with pulsar glitches”. In: *Journal of Physics Conference Series*. Vol. 363. Journal of Physics Conference Series. IOP, 012042, p. 012042. DOI: [10.1088/1742-6596/363/1/012042](https://doi.org/10.1088/1742-6596/363/1/012042). arXiv: [1203.0401](https://arxiv.org/abs/1203.0401) [gr-qc].
- Sarkissian, John, George Hobbs, John Reynolds, Jim Palfreyman, and Steve Olney (Feb. 2019). “Glitch detected in the Vela Pulsar (PSR J0835-4510)”. In: *The Astronomer’s Telegram* 12466, p. 1.
- Sarkissian, John M., Ettore Carretti, and Willem van Straten (Aug. 2011). “The Parkes Pulsar Backends”. In: *Radio Pulsars: An Astrophysical Key to Unlock the Secrets of the Universe*. Ed. by Marta Burgay, Nicolò D’Amico, Paolo Esposito, Alberto Pellizzoni, and Andrea Possenti. Vol. 1357. American Institute of Physics Conference Series. AIP, pp. 351–352. DOI: [10.1063/1.3615155](https://doi.org/10.1063/1.3615155).
- Sarkissian, John M., John E. Reynolds, George Hobbs, and Lisa Harvey-Smith (July 2017). “One Year of Monitoring the Vela Pulsar Using a Phased Array Feed”. In: *Publications of the Astronomical Society of Australia* 34, e027, e027. DOI: [10.1017/pasa.2017.19](https://doi.org/10.1017/pasa.2017.19). arXiv: [1705.08355](https://arxiv.org/abs/1705.08355) [astro-ph.HE].
- Sartore, N., E. Ripamonti, A. Treves, and R. Turolla (Feb. 2010). “Galactic neutron stars. I. Space and velocity distributions in the disk and in the halo”. In: *Astronomy and Astrophysics* 510, A23, A23. DOI: [10.1051/0004-6361/200912222](https://doi.org/10.1051/0004-6361/200912222). arXiv: [0908.3182](https://arxiv.org/abs/0908.3182) [astro-ph.GA].
- Sazhin, M. V. (Feb. 1978). “Opportunities for detecting ultralong gravitational waves”. In: *Soviet Astronomy* 22, pp. 36–38.
- Sazonov, S. et al. (June 2020). “The Galactic LMXB Population and the Galactic Centre Region”. In: *New Astronomy Reviews* 88, 101536, p. 101536. DOI: [10.1016/j.newar.2020.101536](https://doi.org/10.1016/j.newar.2020.101536). arXiv: [2006.05063](https://arxiv.org/abs/2006.05063) [astro-ph.HE].

- Shannon, R. M., L. T. Lentati, M. Kerr, S. Johnston, G. Hobbs, and R. N. Manchester (July 2016). “Characterizing the rotational irregularities of the Vela pulsar from 21 yr of phase-coherent timing”. In: *Monthly Notices of the Royal Astronomical Society* 459.3, pp. 3104–3111. DOI: [10.1093/mnras/stw842](https://doi.org/10.1093/mnras/stw842). arXiv: [1604.02661](https://arxiv.org/abs/1604.02661) [astro-ph.HE].
- Shannon, R. M. et al. (Mar. 2013). “An Asteroid Belt Interpretation for the Timing Variations of the Millisecond Pulsar B1937+21”. In: *The Astrophysical Journal* 766.1, 5, p. 5. DOI: [10.1088/0004-637X/766/1/5](https://doi.org/10.1088/0004-637X/766/1/5). arXiv: [1301.6429](https://arxiv.org/abs/1301.6429) [astro-ph.SR].
- Shannon, R. M. et al. (Sept. 2014). “Limitations in timing precision due to single-pulse shape variability in millisecond pulsars”. In: *Monthly Notices of the Royal Astronomical Society* 443.2, pp. 1463–1481. DOI: [10.1093/mnras/stu1213](https://doi.org/10.1093/mnras/stu1213). arXiv: [1406.4716](https://arxiv.org/abs/1406.4716) [astro-ph.SR].
- Shannon, Ryan M. and James M. Cordes (Dec. 2010). “Assessing the Role of Spin Noise in the Precision Timing of Millisecond Pulsars”. In: *The Astrophysical Journal* 725.2, pp. 1607–1619. DOI: [10.1088/0004-637X/725/2/1607](https://doi.org/10.1088/0004-637X/725/2/1607). arXiv: [1010.4794](https://arxiv.org/abs/1010.4794) [astro-ph.SR].
- Shaw, Benjamin, Mitchell B. Mickaliger, Benjamin W. Stappers, Andrew G. Lyne, Michael J. Keith, Patrick Weltevrede, and Avishek Basu (Sept. 2022). “A very large glitch in the radio pulsar PSR B0740-28”. In: *The Astronomer’s Telegram* 15622, p. 1.
- Shemar, S. L. and A. G. Lyne (Sept. 1996). “Observations of pulsar glitches”. In: *Monthly Notices of the Royal Astronomical Society* 282.2, pp. 677–690. DOI: [10.1093/mnras/282.2.677](https://doi.org/10.1093/mnras/282.2.677).
- Shoemaker, David, Alessio Rocchi, and Shinji Miyoki (2024). *Observing Scenario timeline graphic, post-O3*. LIGO Document G2002127-v26. URL: <https://dcc.ligo.org/LIGO-G2002127-v26/public>.
- Sigg, Daniel (2016). *H1 Calibrated Sensitivity Spectra Oct 24 2015 (Representative for O1)*. LIGO Document LIGO-G1600150. URL: <https://dcc.ligo.org/LIGO-G1600150/public>.
- Singha, Jaikhomba, Avishek Basu, M. A. Krishnakumar, Bhal Chandra Joshi, and P. Arumugam (Aug. 2021a). “A real-time automated glitch detection pipeline at Ooty Radio Telescope”. In: *Monthly Notices of the Royal Astronomical Society* 505.4, pp. 5488–5496. DOI: [10.1093/mnras/stab1640](https://doi.org/10.1093/mnras/stab1640). arXiv: [2106.01942](https://arxiv.org/abs/2106.01942) [astro-ph.IM].
- Singha, Jaikhomba, Bhal Chandra Joshi, Paramasivan Arumugam, and Debades Bandyopadhyay (July 2021b). “Confirmation of recent glitch in PSR J0835-4510(Vela) at 325 MHz using Ooty Radio telescope (ORT)”. in: *The Astronomer’s Telegram* 14812, p. 1.
- Sosa-Fiscella, V. et al. (July 2021). “A new Glitch in the Vela Pulsar (PSR B0833-45/PSR J0835-4510)”. In: *The Astronomer’s Telegram* 14806, p. 1.
- Speagle, Joshua S. (Apr. 2020). “DYNESTY: a dynamic nested sampling package for estimating Bayesian posteriors and evidences”. In: *Monthly Notices of the Royal Astronomical Society* 493.3, pp. 3132–3158. DOI: [10.1093/mnras/staa278](https://doi.org/10.1093/mnras/staa278). arXiv: [1904.02180](https://arxiv.org/abs/1904.02180) [astro-ph.IM].
- Spiewak, R. et al. (Mar. 2018). “PSR J2322-2650 - a low-luminosity millisecond pulsar with a planetary-mass companion”. In: *Monthly Notices of the Royal Astronomical Society* 475.1, pp. 469–477. DOI: [10.1093/mnras/stx3157](https://doi.org/10.1093/mnras/stx3157). arXiv: [1712.04445](https://arxiv.org/abs/1712.04445) [astro-ph.HE].
- Stappers, B. W. et al. (June 2011). “Observing pulsars and fast transients with LOFAR”. in: *Astronomy and Astrophysics* 530, A80, A80. DOI: [10.1051/0004-6361/201116681](https://doi.org/10.1051/0004-6361/201116681). arXiv: [1104.1577](https://arxiv.org/abs/1104.1577) [astro-ph.IM].
- Starobinskiĭ, A. A. (July 1973). “Amplification of waves during reflection from a rotating “black hole””. In: *Soviet Journal of Experimental and Theoretical Physics* 37, p. 28.

- Steltner, B., M. A. Papa, H. B. Eggenstein, R. Prix, M. Bensch, B. Allen, and B. Machenschalk (July 2023). “Deep Einstein@Home All-sky Search for Continuous Gravitational Waves in LIGO O3 Public Data”. In: *The Astrophysical Journal* 952.1, 55, p. 55. DOI: [10.3847/1538-4357/acdad4](https://doi.org/10.3847/1538-4357/acdad4). arXiv: [2303.04109](https://arxiv.org/abs/2303.04109) [gr-qc].
- Stinebring, D. R., M. A. McLaughlin, J. M. Cordes, K. M. Becker, J. E. Espinoza Goodman, M. A. Kramer, J. L. Sheckard, and C. T. Smith (Mar. 2001). “Faint Scattering Around Pulsars: Probing the Interstellar Medium on Solar System Size Scales”. In: *The Astrophysical Journal* 549.1, pp. L97–L100. DOI: [10.1086/319133](https://doi.org/10.1086/319133). arXiv: [astro-ph/0010363](https://arxiv.org/abs/astro-ph/0010363) [astro-ph].
- Sun, L., A. Melatos, S. Suvorova, W. Moran, and R. J. Evans (Feb. 2018). “Hidden Markov model tracking of continuous gravitational waves from young supernova remnants”. In: *Physical Review D* 97.4, 043013, p. 043013. DOI: [10.1103/PhysRevD.97.043013](https://doi.org/10.1103/PhysRevD.97.043013). arXiv: [1710.00460](https://arxiv.org/abs/1710.00460) [astro-ph.IM].
- Sun, Ling, Richard Brito, and Maximiliano Isi (Mar. 2020a). “Search for ultralight bosons in Cygnus X-1 with Advanced LIGO”. in: *Physical Review D* 101.6, 063020, p. 063020. DOI: [10.1103/PhysRevD.101.063020](https://doi.org/10.1103/PhysRevD.101.063020). arXiv: [1909.11267](https://arxiv.org/abs/1909.11267) [gr-qc].
- Sun, Ling and Andrew Melatos (June 2019). “Application of hidden Markov model tracking to the search for long-duration transient gravitational waves from the remnant of the binary neutron star merger GW170817”. In: *Physical Review D* 99.12, 123003, p. 123003. DOI: [10.1103/PhysRevD.99.123003](https://doi.org/10.1103/PhysRevD.99.123003). arXiv: [1810.03577](https://arxiv.org/abs/1810.03577) [astro-ph.IM].
- Sun, Ling, Andrew Melatos, and Paul D. Lasky (June 2019). “Tracking continuous gravitational waves from a neutron star at once and twice the spin frequency with a hidden Markov model”. In: *Physical Review D* 99.12, 123010, p. 123010. DOI: [10.1103/PhysRevD.99.123010](https://doi.org/10.1103/PhysRevD.99.123010). arXiv: [1903.03866](https://arxiv.org/abs/1903.03866) [astro-ph.IM].
- Sun, Ling et al. (Nov. 2020b). “Characterization of systematic error in Advanced LIGO calibration”. In: *Classical and Quantum Gravity* 37.22, 225008, p. 225008. DOI: [10.1088/1361-6382/abb14e](https://doi.org/10.1088/1361-6382/abb14e). arXiv: [2005.02531](https://arxiv.org/abs/2005.02531) [astro-ph.IM].
- Sun, Ling et al. (June 2021). “Characterization of systematic error in Advanced LIGO calibration in the second half of O3”. In: *arXiv e-prints*, arXiv:2107.00129, arXiv:2107.00129. DOI: [10.48550/arXiv.2107.00129](https://doi.org/10.48550/arXiv.2107.00129). arXiv: [2107.00129](https://arxiv.org/abs/2107.00129) [astro-ph.IM].
- Suvorov, A. G. and A. Melatos (Mar. 2019). “Relaxation by thermal conduction of a magnetically confined mountain on an accreting neutron star”. In: *Monthly Notices of the Royal Astronomical Society* 484.1, pp. 1079–1099. DOI: [10.1093/mnras/sty3518](https://doi.org/10.1093/mnras/sty3518). arXiv: [1812.10029](https://arxiv.org/abs/1812.10029) [astro-ph.HE].
- Suvorova, S., P. Clearwater, A. Melatos, L. Sun, W. Moran, and R. J. Evans (Nov. 2017). “Hidden Markov model tracking of continuous gravitational waves from a binary neutron star with wandering spin. II. Binary orbital phase tracking”. In: *Physical Review D* 96.10, 102006, p. 102006. DOI: [10.1103/PhysRevD.96.102006](https://doi.org/10.1103/PhysRevD.96.102006). arXiv: [1710.07092](https://arxiv.org/abs/1710.07092) [astro-ph.IM].
- Suvorova, S., L. Sun, A. Melatos, W. Moran, and R. J. Evans (June 2016). “Hidden Markov model tracking of continuous gravitational waves from a neutron star with wandering spin”. In: *Physical Review D* 93.12, 123009, p. 123009. DOI: [10.1103/PhysRevD.93.123009](https://doi.org/10.1103/PhysRevD.93.123009). arXiv: [1606.02412](https://arxiv.org/abs/1606.02412) [astro-ph.IM].

- Suvorova, Sofia, Andrew Melatos, Rob J. Evans, William Moran, Patrick Clearwater, and Ling Sun (2018). “Phase-Continuous Frequency Line Track-Before-Detect of a Tone With Slow Frequency Variation”. In: *IEEE Transactions on Signal Processing* 66.24, pp. 6434–6442. DOI: [10.1109/TSP.2018.2877176](https://doi.org/10.1109/TSP.2018.2877176).
- Tan, C. M. et al. (Oct. 2018). “LOFAR Discovery of a 23.5 s Radio Pulsar”. In: *The Astrophysical Journal* 866.1, 54, p. 54. DOI: [10.3847/1538-4357/aade88](https://doi.org/10.3847/1538-4357/aade88). arXiv: [1809.00965](https://arxiv.org/abs/1809.00965) [astro-ph.HE].
- Tauris, T. M., D. Sanyal, S. C. Yoon, and N. Langer (Oct. 2013). “Evolution towards and beyond accretion-induced collapse of massive white dwarfs and formation of millisecond pulsars”. In: *Astronomy and Astrophysics* 558, A39, A39. DOI: [10.1051/0004-6361/201321662](https://doi.org/10.1051/0004-6361/201321662). arXiv: [1308.4887](https://arxiv.org/abs/1308.4887) [astro-ph.SR].
- Taylor, J. H. (Oct. 1992). “Pulsar Timing and Relativistic Gravity”. In: *Philosophical Transactions of the Royal Society of London Series A* 341.1660, pp. 117–134. DOI: [10.1098/rsta.1992.0088](https://doi.org/10.1098/rsta.1992.0088).
- Taylor, J. H., L. A. Fowler, and P. M. McCulloch (Feb. 1979). “Measurements of general relativistic effects in the binary pulsar PSR1913 + 16”. In: *Nature* 277.5696, pp. 437–440. DOI: [10.1038/277437a0](https://doi.org/10.1038/277437a0).
- Taylor Joseph H., Jr. (July 1991). “Millisecond pulsars: nature’s most stable clocks.” In: *IEEE Proceedings* 79, pp. 1054–1062.
- Taylor, Stephen R. (May 2021). “The Nanohertz Gravitational Wave Astronomer”. In: *arXiv e-prints*, arXiv:2105.13270, arXiv:2105.13270. DOI: [10.48550/arXiv.2105.13270](https://doi.org/10.48550/arXiv.2105.13270). arXiv: [2105.13270](https://arxiv.org/abs/2105.13270) [astro-ph.HE].
- Tenorio, Rodrigo, David Keitel, and Alicia M. Sintes (Oct. 2021). “Application of a hierarchical MCMC follow-up to Advanced LIGO continuous gravitational-wave candidates”. In: *Physical Review D* 104.8, 084012, p. 084012. DOI: [10.1103/PhysRevD.104.084012](https://doi.org/10.1103/PhysRevD.104.084012). arXiv: [2105.13860](https://arxiv.org/abs/2105.13860) [gr-qc].
- Thompson, Christopher and Robert C. Duncan (May 1993). “Neutron Star Dynamos and the Origins of Pulsar Magnetism”. In: *The Astrophysical Journal* 408, p. 194. DOI: [10.1086/172580](https://doi.org/10.1086/172580).
- Thong, K. H., A. Melatos, and L. V. Drummond (June 2023). “Stability of interlinked neutron vortex and proton flux-tube arrays in a neutron star - III. Proton feedback”. In: *Monthly Notices of the Royal Astronomical Society* 521.4, pp. 5724–5737. DOI: [10.1093/mnras/stad927](https://doi.org/10.1093/mnras/stad927). arXiv: [2305.04482](https://arxiv.org/abs/2305.04482) [astro-ph.HE].
- Thorne, Kip S. (Apr. 1980). “Gravitational-wave research: Current status and future prospects”. In: *Reviews of Modern Physics* 52.2, pp. 285–299. DOI: [10.1103/RevModPhys.52.285](https://doi.org/10.1103/RevModPhys.52.285).
- Tiwari, V. et al. (Aug. 2015). “Regression of environmental noise in LIGO data”. In: *Classical and Quantum Gravity* 32.16, 165014, p. 165014. DOI: [10.1088/0264-9381/32/16/165014](https://doi.org/10.1088/0264-9381/32/16/165014). arXiv: [1503.07476](https://arxiv.org/abs/1503.07476) [gr-qc].
- Turk, P. J. and D. R. Lorimer (Dec. 2013). “An empirical Bayesian analysis applied to the globular cluster pulsar population”. In: *Monthly Notices of the Royal Astronomical Society* 436.4, pp. 3720–3726. DOI: [10.1093/mnras/stt1850](https://doi.org/10.1093/mnras/stt1850). arXiv: [1309.7317](https://arxiv.org/abs/1309.7317) [astro-ph.GA].

- Urama, J. O., B. Link, and J. M. Weisberg (July 2006). “A strong correlation in radio pulsars with implications for torque variations”. In: *Monthly Notices of the Royal Astronomical Society* 370.1, pp. L76–L79. DOI: [10.1111/j.1745-3933.2006.00192.x](https://doi.org/10.1111/j.1745-3933.2006.00192.x). arXiv: [astro-ph/0605380](https://arxiv.org/abs/astro-ph/0605380) [astro-ph].
- Ushomirsky, Greg, Curt Cutler, and Lars Bildsten (Dec. 2000). “Deformations of accreting neutron star crusts and gravitational wave emission”. In: *Monthly Notices of the Royal Astronomical Society* 319.3, pp. 902–932. DOI: [10.1046/j.1365-8711.2000.03938.x](https://doi.org/10.1046/j.1365-8711.2000.03938.x). arXiv: [astro-ph/0001136](https://arxiv.org/abs/astro-ph/0001136) [astro-ph].
- Vallisneri, Michele (Feb. 2020). *libstempo: Python wrapper for Tempo2*. Astrophysics Source Code Library, record ascl:2002.017.
- van der Horst, A. J. et al. (Apr. 2012). “SGR J1550-5418 Bursts Detected with the Fermi Gamma-Ray Burst Monitor during its Most Prolific Activity”. In: *The Astrophysical Journal* 749.2, 122, p. 122. DOI: [10.1088/0004-637X/749/2/122](https://doi.org/10.1088/0004-637X/749/2/122). arXiv: [1202.3157](https://arxiv.org/abs/1202.3157) [astro-ph.HE].
- van Haasteren, Rutger and Yuri Levin (Jan. 2013). “Understanding and analysing time-correlated stochastic signals in pulsar timing”. In: *Monthly Notices of the Royal Astronomical Society* 428.2, pp. 1147–1159. DOI: [10.1093/mnras/sts097](https://doi.org/10.1093/mnras/sts097). arXiv: [1202.5932](https://arxiv.org/abs/1202.5932) [astro-ph.IM].
- van Haasteren, Rutger, Yuri Levin, Patrick McDonald, and Tingting Lu (May 2009). “On measuring the gravitational-wave background using Pulsar Timing Arrays”. In: *Monthly Notices of the Royal Astronomical Society* 395.2, pp. 1005–1014. DOI: [10.1111/j.1365-2966.2009.14590.x](https://doi.org/10.1111/j.1365-2966.2009.14590.x). arXiv: [0809.0791](https://arxiv.org/abs/0809.0791) [astro-ph].
- van Straten, W. and M. Bailes (Jan. 2011). “DSPSR: Digital Signal Processing Software for Pulsar Astronomy”. In: *Publications of the Astronomical Society of Australia* 28.1, pp. 1–14. DOI: [10.1071/AS10021](https://doi.org/10.1071/AS10021). arXiv: [1008.3973](https://arxiv.org/abs/1008.3973) [astro-ph.IM].
- Vargas, Andrés F. and Andrew Melatos (July 2023a). “Apparent dispersion in pulsar braking index measurements caused by timing noise”. In: *Monthly Notices of the Royal Astronomical Society* 522.4, pp. 4880–4893. DOI: [10.1093/mnras/stad1301](https://doi.org/10.1093/mnras/stad1301). arXiv: [2305.09079](https://arxiv.org/abs/2305.09079) [astro-ph.HE].
- (Mar. 2023b). “Search for continuous gravitational waves from PSR J 0437 -4715 with a hidden Markov model in O3 LIGO data”. In: *Physical Review D* 107.6, 064062, p. 064062. DOI: [10.1103/PhysRevD.107.064062](https://doi.org/10.1103/PhysRevD.107.064062). arXiv: [2208.03932](https://arxiv.org/abs/2208.03932) [gr-qc].
- Vargas, F.L., R.D.R. Fagundes, and D.B. Junior (2001). “A FPGA-based Viterbi algorithm implementation for speech recognition systems”. In: *2001 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No.01CH37221)*. Vol. 2, 1217–1220 vol.2. DOI: [10.1109/ICASSP.2001.941143](https://doi.org/10.1109/ICASSP.2001.941143).
- Venkatraman Krishnan, V. et al. (Mar. 2020). “The UTMOST survey for magnetars, intermittent pulsars, RRATs, and FRBs - I. System description and overview”. In: *Monthly Notices of the Royal Astronomical Society* 492.4, pp. 4752–4767. DOI: [10.1093/mnras/staa111](https://doi.org/10.1093/mnras/staa111). arXiv: [1905.02415](https://arxiv.org/abs/1905.02415) [astro-ph.IM].
- Verbiest, Joris P. W. and G. M. Shaifullah (July 2018). “Measurement uncertainty in pulsar timing array experiments”. In: *Classical and Quantum Gravity* 35.13, 133001, p. 133001. DOI: [10.1088/1361-6382/aac412](https://doi.org/10.1088/1361-6382/aac412).
- Verbunt, F. and P. Hut (Jan. 1987). “The Globular Cluster Population of X-Ray Binaries”. In: *The Origin and Evolution of Neutron Stars*. Ed. by D. J. Helfand and J. H. Huang. Vol. 125. IAU Symposium, p. 187.

- Verbunt, Frank and Paulo C. C. Freire (Jan. 2014). “On the disruption of pulsar and X-ray binaries in globular clusters”. In: *Astronomy and Astrophysics* 561, A11, A11. DOI: [10.1051/0004-6361/201321177](https://doi.org/10.1051/0004-6361/201321177). arXiv: [1310.4669](https://arxiv.org/abs/1310.4669) [astro-ph.SR].
- Vigelius, M. and A. Melatos (June 2009). “Resistive relaxation of a magnetically confined mountain on an accreting neutron star”. In: *Monthly Notices of the Royal Astronomical Society* 395.4, pp. 1985–1998. DOI: [10.1111/j.1365-2966.2009.14698.x](https://doi.org/10.1111/j.1365-2966.2009.14698.x). arXiv: [0902.4484](https://arxiv.org/abs/0902.4484) [astro-ph.HE].
- Viterbi, A. (1967). “Error bounds for convolutional codes and an asymptotically optimum decoding algorithm”. In: *IEEE Transactions on Information Theory* 13.2, pp. 260–269. DOI: [10.1109/TIT.1967.1054010](https://doi.org/10.1109/TIT.1967.1054010).
- Vleschower, L. et al. (June 2022). “Discoveries and timing of pulsars in NGC 6440”. In: *Monthly Notices of the Royal Astronomical Society* 513.1, pp. 1386–1399. DOI: [10.1093/mnras/stac921](https://doi.org/10.1093/mnras/stac921). arXiv: [2204.00086](https://arxiv.org/abs/2204.00086) [astro-ph.HE].
- Vurgun, Eda et al. (Dec. 2022). “The Neutron Star Population in M28: A Joint Chandra/GBT Look at Pulsar Paradise”. In: *The Astrophysical Journal* 941.1, 76, p. 76. DOI: [10.3847/1538-4357/ac9ea0](https://doi.org/10.3847/1538-4357/ac9ea0). arXiv: [2211.01067](https://arxiv.org/abs/2211.01067) [astro-ph.HE].
- Wagoner, R. V. (Mar. 1984). “Gravitational radiation from accreting neutron stars”. In: *The Astrophysical Journal* 278, pp. 345–348. DOI: [10.1086/161798](https://doi.org/10.1086/161798).
- Wang, J., J. P. W. Verbiest, G. M. Shaifullah, I. Cognard, L. Guillemot, G. H. Janssen, M. B. Mickaliger, A. Possenti, and G. Theureau (July 2024). “Improving pulsar timing precision through superior time-of-arrival creation”. In: *Astronomy and Astrophysics* 687, A154, A154. DOI: [10.1051/0004-6361/202449366](https://doi.org/10.1051/0004-6361/202449366). arXiv: [2405.08629](https://arxiv.org/abs/2405.08629) [astro-ph.IM].
- Wang, N., R. N. Manchester, and S. Johnston (May 2007). “Pulsar nulling and mode changing”. In: *Monthly Notices of the Royal Astronomical Society* 377.3, pp. 1383–1392. DOI: [10.1111/j.1365-2966.2007.11703.x](https://doi.org/10.1111/j.1365-2966.2007.11703.x). arXiv: [astro-ph/0703241](https://arxiv.org/abs/astro-ph/0703241) [astro-ph].
- Wang, N., R. N. Manchester, R. T. Pace, M. Bailes, V. M. Kaspi, B. W. Stappers, and A. G. Lyne (Oct. 2000). “Glitches in southern pulsars”. In: *Monthly Notices of the Royal Astronomical Society* 317.4, pp. 843–860. DOI: [10.1046/j.1365-8711.2000.03713.x](https://doi.org/10.1046/j.1365-8711.2000.03713.x). arXiv: [astro-ph/0005561](https://arxiv.org/abs/astro-ph/0005561) [astro-ph].
- Wang, Zhongxiang, Deepto Chakrabarty, and David L. Kaplan (Apr. 2006). “A debris disk around an isolated young neutron star”. In: *Nature* 440.7085, pp. 772–775. DOI: [10.1038/nature04669](https://doi.org/10.1038/nature04669). arXiv: [astro-ph/0604076](https://arxiv.org/abs/astro-ph/0604076) [astro-ph].
- Warszawski, L. and A. Melatos (Jan. 2013). “Knock-on processes in superfluid vortex avalanches and pulsar glitch statistics”. In: *Monthly Notices of the Royal Astronomical Society* 428.3, pp. 1911–1926. DOI: [10.1093/mnras/sts108](https://doi.org/10.1093/mnras/sts108). arXiv: [1210.2203](https://arxiv.org/abs/1210.2203) [astro-ph.HE].
- Warszawski, L., A. Melatos, and N. G. Berloff (Mar. 2012). “Unpinning triggers for superfluid vortex avalanches”. In: *Physical Review B* 85.10, 104503, p. 104503. DOI: [10.1103/PhysRevB.85.104503](https://doi.org/10.1103/PhysRevB.85.104503). arXiv: [1203.5133](https://arxiv.org/abs/1203.5133) [cond-mat.other].
- Weber, J. (June 1969). “Evidence for Discovery of Gravitational Radiation”. In: *Physical Review Letters* 22.24, pp. 1320–1324. DOI: [10.1103/PhysRevLett.22.1320](https://doi.org/10.1103/PhysRevLett.22.1320).
- (July 1970). “Anisotropy and Polarization in the Gravitational-Radiation Experiments”. In: *Physical Review Letters* 25.3, pp. 180–184. DOI: [10.1103/PhysRevLett.25.180](https://doi.org/10.1103/PhysRevLett.25.180).
- Weber, Joseph and John A. Wheeler (July 1957). “Reality of the Cylindrical Gravitational Waves of Einstein and Rosen”. In: *Reviews of Modern Physics* 29.3, pp. 509–515. DOI: [10.1103/RevModPhys.29.509](https://doi.org/10.1103/RevModPhys.29.509).

- Weisberg, J. M., D. J. Nice, and J. H. Taylor (Oct. 2010). “Timing Measurements of the Relativistic Binary Pulsar PSR B1913+16”. In: *The Astrophysical Journal* 722.2, pp. 1030–1034. DOI: [10.1088/0004-637X/722/2/1030](https://doi.org/10.1088/0004-637X/722/2/1030). arXiv: [1011.0718](https://arxiv.org/abs/1011.0718) [astro-ph.GA].
- Weltevrede, Patrick, Simon Johnston, and Cristóbal M. Espinoza (Mar. 2011). “The glitch-induced identity changes of PSR J1119-6127”. In: *Monthly Notices of the Royal Astronomical Society* 411.3, pp. 1917–1934. DOI: [10.1111/j.1365-2966.2010.17821.x](https://doi.org/10.1111/j.1365-2966.2010.17821.x). arXiv: [1010.0857](https://arxiv.org/abs/1010.0857) [astro-ph.SR].
- Wette, K., S. Walsh, R. Prix, and M. A. Papa (June 2018). “Implementing a semicoherent search for continuous gravitational waves using optimally constructed template banks”. In: *Physical Review D* 97.12, 123016, p. 123016. DOI: [10.1103/PhysRevD.97.123016](https://doi.org/10.1103/PhysRevD.97.123016). arXiv: [1804.03392](https://arxiv.org/abs/1804.03392) [astro-ph.IM].
- Wette, K. et al. (Dec. 2008). “Searching for gravitational waves from Cassiopeia A with LIGO”. In: *Classical and Quantum Gravity* 25.23, 235011, p. 235011. DOI: [10.1088/0264-9381/25/23/235011](https://doi.org/10.1088/0264-9381/25/23/235011). arXiv: [0802.3332](https://arxiv.org/abs/0802.3332) [gr-qc].
- Wette, Karl (Feb. 2012). “Estimating the sensitivity of wide-parameter-space searches for gravitational-wave pulsars”. In: *Physical Review D* 85.4, 042003, p. 042003. DOI: [10.1103/PhysRevD.85.042003](https://doi.org/10.1103/PhysRevD.85.042003). arXiv: [1111.5650](https://arxiv.org/abs/1111.5650) [gr-qc].
- (Oct. 2015). “Parameter-space metric for all-sky semicoherent searches for gravitational-wave pulsars”. In: *Physical Review D* 92.8, 082003, p. 082003. DOI: [10.1103/PhysRevD.92.082003](https://doi.org/10.1103/PhysRevD.92.082003). arXiv: [1508.02372](https://arxiv.org/abs/1508.02372) [gr-qc].
 - (Nov. 2023). “Searches for continuous gravitational waves from neutron stars: A twenty-year retrospective”. In: *Astroparticle Physics* 153, 102880, p. 102880. DOI: [10.1016/j.astropartphys.2023.102880](https://doi.org/10.1016/j.astropartphys.2023.102880). arXiv: [2305.07106](https://arxiv.org/abs/2305.07106) [gr-qc].
- Wette, Karl, Liam Dunn, Patrick Clearwater, and Andrew Melatos (Apr. 2021). “Deep exploration for continuous gravitational waves at 171-172 Hz in LIGO second observing run data”. In: *Physical Review D* 103.8, 083020, p. 083020. DOI: [10.1103/PhysRevD.103.083020](https://doi.org/10.1103/PhysRevD.103.083020). arXiv: [2103.12976](https://arxiv.org/abs/2103.12976) [gr-qc].
- Wette, Karl William (Jan. 2009). “Gravitational waves from accreting neutron stars and Cassiopeia A”. PhD thesis. Australian National University, Canberra.
- Whelan, John T., Reinhard Prix, Curt J. Cutler, and Joshua L. Willis (Mar. 2014). “New coordinates for the amplitude parameter space of continuous gravitational waves”. In: *Classical and Quantum Gravity* 31.6, 065002, p. 065002. DOI: [10.1088/0264-9381/31/6/065002](https://doi.org/10.1088/0264-9381/31/6/065002). arXiv: [1311.0065](https://arxiv.org/abs/1311.0065) [gr-qc].
- Whelan, John T., Santosh Sundaesan, Yuanhao Zhang, and Prabath Peiris (May 2015). “Model-based cross-correlation search for gravitational waves from Scorpius X-1”. In: *Physical Review D* 91.10, 102005, p. 102005. DOI: [10.1103/PhysRevD.91.102005](https://doi.org/10.1103/PhysRevD.91.102005). arXiv: [1504.05890](https://arxiv.org/abs/1504.05890) [gr-qc].
- White, Jack, Karel Adámek, Jayanta Roy, Scott Ransom, and Wesley Armour (June 2024). “Pulscan: Binary pulsar detection using unmatched filters on NVIDIA GPUs”. In: *arXiv e-prints*, arXiv:2406.15186, arXiv:2406.15186. DOI: [10.48550/arXiv.2406.15186](https://doi.org/10.48550/arXiv.2406.15186). arXiv: [2406.15186](https://arxiv.org/abs/2406.15186) [astro-ph.IM].
- Wolszczan, A. and D. A. Frail (Jan. 1992). “A planetary system around the millisecond pulsar PSR1257 + 12”. In: *Nature* 355.6356, pp. 145–147. DOI: [10.1038/355145a0](https://doi.org/10.1038/355145a0).
- Wong, Tony, D. C. Backer, and A. G. Lyne (Feb. 2001). “Observations of a Series of Six Recent Glitches in the Crab Pulsar”. In: *The Astrophysical Journal* 548.1, pp. 447–459. DOI: [10.1086/318657](https://doi.org/10.1086/318657). arXiv: [astro-ph/0010010](https://arxiv.org/abs/astro-ph/0010010) [astro-ph].

- Wood, Toby S. and Vanessa Graber (Apr. 2022). “Superconducting Phases in Neutron Star Cores”. In: *Universe* 8.4, 228, p. 228. DOI: [10.3390/universe8040228](https://doi.org/10.3390/universe8040228).
- Xu, H. et al. (May 2021). “A sustained pulse shape change in PSR J1713+0747 possibly associated with timing and DM events”. In: *The Astronomer’s Telegram* 14642, p. 1.
- Xu, Heng et al. (July 2023). “Searching for the Nano-Hertz Stochastic Gravitational Wave Background with the Chinese Pulsar Timing Array Data Release I”. in: *Research in Astronomy and Astrophysics* 23.7, 075024, p. 075024. DOI: [10.1088/1674-4527/acdfa5](https://doi.org/10.1088/1674-4527/acdfa5). arXiv: [2306.16216](https://arxiv.org/abs/2306.16216) [astro-ph.HE].
- Ye, Claire S. and Giacomo Fragione (Dec. 2022). “Millisecond Pulsars in Dense Star Clusters: Evolution, Scaling Relations, and the Galactic-center Gamma-Ray Excess”. In: *The Astrophysical Journal* 940.2, 162, p. 162. DOI: [10.3847/1538-4357/ac9cd0](https://doi.org/10.3847/1538-4357/ac9cd0). arXiv: [2207.03504](https://arxiv.org/abs/2207.03504) [astro-ph.HE].
- Ye, Claire S., Kyle Kremer, Sourav Chatterjee, Carl L. Rodriguez, and Frederic A. Rasio (June 2019). “Millisecond Pulsars and Black Holes in Globular Clusters”. In: *The Astrophysical Journal* 877.2, 122, p. 122. DOI: [10.3847/1538-4357/ab1b21](https://doi.org/10.3847/1538-4357/ab1b21). arXiv: [1902.05963](https://arxiv.org/abs/1902.05963) [astro-ph.HE].
- Yin, Dejiang et al. (July 2024). “FAST Discovery of Eight Isolated Millisecond Pulsars in NGC 6517”. In: *The Astrophysical Journal* 969.1, L7, p. L7. DOI: [10.3847/2041-8213/ad534e](https://doi.org/10.3847/2041-8213/ad534e). arXiv: [2405.18228](https://arxiv.org/abs/2405.18228) [astro-ph.HE].
- Yu, M. and Q. J. Liu (July 2017). “On the detection probability of neutron star glitches”. In: *Monthly Notices of the Royal Astronomical Society* 468.3, pp. 3031–3041. DOI: [10.1093/mnras/stx702](https://doi.org/10.1093/mnras/stx702). arXiv: [1704.05983](https://arxiv.org/abs/1704.05983) [astro-ph.HE].
- Yu, M. et al. (Feb. 2013). “Detection of 107 glitches in 36 southern pulsars”. In: *Monthly Notices of the Royal Astronomical Society* 429.1, pp. 688–724. DOI: [10.1093/mnras/sts366](https://doi.org/10.1093/mnras/sts366). arXiv: [1211.2035](https://arxiv.org/abs/1211.2035) [astro-ph.HE].
- Yu, Meng (Nov. 2018). “Harmonic Summing Improves Pulsar Detection Sensitivity: A Probability Analysis”. In: *The Astrophysical Journal* 868.1, 8, p. 8. DOI: [10.3847/1538-4357/aae51a](https://doi.org/10.3847/1538-4357/aae51a). arXiv: [1811.07111](https://arxiv.org/abs/1811.07111) [astro-ph.IM].
- Yuan, J. P., N. Wang, R. N. Manchester, and Z. Y. Liu (May 2010). “29 glitches detected at Urumqi Observatory”. In: *Monthly Notices of the Royal Astronomical Society* 404.1, pp. 289–304. DOI: [10.1111/j.1365-2966.2010.16272.x](https://doi.org/10.1111/j.1365-2966.2010.16272.x). arXiv: [1001.1471](https://arxiv.org/abs/1001.1471) [astro-ph.GA].
- Zel’Dovich, Ya. B. (Aug. 1971). “Generation of Waves by a Rotating Body”. In: *Soviet Journal of Experimental and Theoretical Physics Letters* 14, p. 180.
- Zhang, Bing, Alice K. Harding, and Alexander G. Muslimov (Mar. 2000). “Radio Pulsar Death Line Revisited: Is PSR J2144-3933 Anomalous?” In: *The Astrophysical Journal* 531.2, pp. L135–L138. DOI: [10.1086/312542](https://doi.org/10.1086/312542). arXiv: [astro-ph/0001341](https://arxiv.org/abs/astro-ph/0001341) [astro-ph].
- Zhang, Pengfei, Yi Xing, and Zhongxiang Wang (Aug. 2022). “Discovery of γ -Ray Pulsations from PSR J1835-3259B in the Globular Cluster NGC 6652”. In: *The Astrophysical Journal* 935.2, L36, p. L36. DOI: [10.3847/2041-8213/ac88bf](https://doi.org/10.3847/2041-8213/ac88bf). arXiv: [2206.13667](https://arxiv.org/abs/2206.13667) [astro-ph.HE].
- Zheng, Jiao, Pengfei Zhang, and Li Zhang (Jan. 2024). “A Possible γ -Ray Pulsation from PSR J1740-5340B in the Globular Cluster NGC 6397”. In: *Research in Astronomy and Astrophysics* 24.1, 015023, p. 015023. DOI: [10.1088/1674-4527/ad1366](https://doi.org/10.1088/1674-4527/ad1366).
- Zubieta, E., S. B. Araujo Furlan, S. del Palacio, F. Garcia, G. Gancio, C. O. Lousto, J. A. Combi, and L. Combi (Dec. 2022a). “Detection of a new glitch in the pulsar PSR J1740-3015”. In: *The Astronomer’s Telegram* 15838, p. 1.

- Zubieta, E., F. Garcia, S. del Palacio, S. B. Araujo Furlan, G. Gancio, C. O. Lousto, J. A. Combi, and C. M. Espinoza (June 2024). “Timing irregularities and glitches from the pulsar monitoring campaign at IAR”. in: *arXiv e-prints*, arXiv:2406.17099, arXiv:2406.17099. DOI: [10.48550/arXiv.2406.17099](https://doi.org/10.48550/arXiv.2406.17099). arXiv: [2406.17099](https://arxiv.org/abs/2406.17099) [astro-ph.HE].
- Zubieta, E. et al. (Sept. 2022b). “Confirmation of glitch event observed in PSR J0742-2822 at the Argentine Institute of Radioastronomy (IAR)”. in: *The Astronomer’s Telegram* 15638, p. 1.
- Zubieta, Ezequiel et al. (May 2023). “First results of the glitching pulsar monitoring programme at the Argentine Institute of Radioastronomy”. In: *Monthly Notices of the Royal Astronomical Society* 521.3, pp. 4504–4521. DOI: [10.1093/mnras/stad723](https://doi.org/10.1093/mnras/stad723). arXiv: [2210.03770](https://arxiv.org/abs/2210.03770) [astro-ph.HE].
- Zweizig, John and Keith Riles (2021). *Information on self-gating of $h(t)$ used in O3 continuous-wave and stochastic searches*. LIGO Document T2000384. URL: <https://dcc.ligo.org/T2000384/public>.

This thesis has been typeset using \LaTeX . The text face is Linux Libertine, designed by Philipp H. Poll. The mathematics is set in KP Fonts Serif, based on URW Palladio. URLs and other identifiers are set in Calibri. The opening quotation is “On Exactitude in Science” by Jorge Luis Borges, translated by Andrew Hurley. The look & feel of this thesis is closely based on the Dissertate template, which has been released under an AGPL license and is available from github.com/suchow/Dissertate.