



HAL
open science

A Gaussian-processes approach to fitting for time-variable spherical solar wind in pulsar timing data

Iuliana C Nițu, Michael J Keith, Caterina Tiburzi, Marcus Brüggen, David J Champion, Siyuan Chen, Ismaël Cognard, Gregory Desvignes, Ralf-Jürgen Dettmar, Jean-Mathias Grießmeier, et al.

► To cite this version:

Iuliana C Nițu, Michael J Keith, Caterina Tiburzi, Marcus Brüggen, David J Champion, et al.. A Gaussian-processes approach to fitting for time-variable spherical solar wind in pulsar timing data. Monthly Notices of the Royal Astronomical Society, 2024, 528 (2), pp.3304-3319. 10.1093/mnras/stae220 . hal-04446854

HAL Id: hal-04446854

<https://hal.science/hal-04446854>

Submitted on 21 Apr 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

A Gaussian-processes approach to fitting for time-variable spherical solar wind in pulsar timing data

Iuliana C. Nițu¹,¹★ Michael J. Keith¹,¹★ Caterina Tiburzi,² Marcus Brüggen³, David J. Champion⁴, Siyuan Chen⁵, Ismaël Cognard^{6,7}, Gregory Desvignes⁴, Ralf-Jürgen Dettmar,⁸ Jean-Mathias Grießmeier^{6,7}, Lucas Guillemot^{6,7}, Yanjun Guo,⁴ Matthias Hoelt⁹, Huanchen Hu⁴, Jiwoong Jang⁴, Gemma H. Janssen^{10,11}, Jędrzej Jawor⁴, Ramesh Karuppusamy⁴, Evan F. Keane¹², Michael Kramer,⁴ Jörn Künsemöller,¹³ Kristen Lackeos⁴, Kuo Liu^{14,4}, Robert A. Main,⁴ James W. McKee^{15,16}, Nataliya K. Porayko,⁴ Golam M. Shaifullah^{2,17,18}, Gilles Theureau^{6,7,19} and Christian Vocks²⁰

Affiliations are listed at the end of the paper

Accepted 2024 January 17. Received 2024 January 16; in original form 2023 November 9

ABSTRACT

Propagation effects are one of the main sources of noise in high-precision pulsar timing. For pulsars below an ecliptic latitude of 5° , the ionized plasma in the solar wind can introduce dispersive delays of order $100 \mu\text{s}$ around solar conjunction at an observing frequency of 300 MHz. A common approach to mitigate this assumes a spherical solar wind with a time-constant amplitude. However, this has been shown to be insufficient to describe the solar wind. We present a linear, Gaussian-process piecewise Bayesian approach to fit a spherical solar wind of time-variable amplitude, which has been implemented in the pulsar software `RUN_ENTERPRISE`. Through simulations, we find that the current EPTA+InPTA data combination is not sensitive to such variations; however, solar wind variations will become important in the near future with the addition of new InPTA data and data collected with the low-frequency LOFAR telescope. We also compare our results for different high-precision timing data sets (EPTA+InPTA, PPTA, and LOFAR) of 3 ms pulsars (J0030+0451, J1022+1001, J2145–0450), and find that the solar-wind amplitudes are generally consistent for any individual pulsar, but they can vary from pulsar to pulsar. Finally, we compare our results with those of an independent method on the same LOFAR data of the three millisecond pulsars. We find that differences between the results of the two methods can be mainly attributed to the modelling of dispersion variations in the interstellar medium, rather than the solar wind modelling.

Key words: methods: data analysis – solar wind – pulsars: general – pulsars: individual: PSR J0030+0451, PSR J1022+1001, PSR J2145–0450.

1 INTRODUCTION

Pulsar timing consists of recording the times of arrival (ToAs) of highly stable pulses emitted by pulsars, which are then compared with predictions from long-term models. These models take into account the pulsar’s behaviour, as well as factors such as astrometric effects (e.g. spin frequency and its derivatives, position, etc.), binary companions, or the dispersive delays induced by the ionized medium through which the pulsar signal propagates (for more details see e.g. Edwards, Hobbs & Manchester 2006). The left-over signal after subtracting the model from the observed ToAs is often referred to as ‘timing residuals’, or just ‘residuals’, and is expected to be only white noise if the model is optimal.

Due to the high stability of pulsar rotation and particularly of the population of recycled millisecond pulsars (MSPs; Backer et al. 1982), high-precision pulsar timing constitutes a great tool for a large variety of scientific investigations (see e.g. Manchester 2017). Huge efforts are currently being concentrated on gravitational wave searches using decades of observations of a large sample of MSPs with multiple telescopes, referred to as Pulsar Timing Arrays (PTAs; e.g. Tiburzi 2018). PTA experiments are expected to be primarily sensitive to the gravitational wave background (GWB) in the nanohertz-frequency regime, most likely originating from supermassive black hole binary mergers (e.g. Burke-Spolaor et al. 2019). Three major collaborations have historically been involved in the search for the GWB, namely the European Pulsar Timing Array (EPTA; Desvignes et al. 2016), the Parkes Pulsar Timing Array (PPTA; Manchester et al. 2013), and the North American Nanohertz Observatory for Gravitational Waves (NANOGrav; Demorest et al. 2013); these are also the founding members of the International Pulsar Timing Array (IPTA; Verbiest et al. 2016) collaboration, having recently

* E-mail: iuliana-camelia.nitu@manchester.ac.uk (ICN);
michael.keith@manchester.ac.uk (MJK)

been joined by the Indian Pulsar Timing Array (InPTA; Joshi et al. 2022). Recently, the three aforementioned collaborations (with the InPTA working along the EPTA) all coherently reported marginal evidence for a GWB signal (Agazie et al. 2023; Reardon et al. 2023; EPTA Collaboration et al. 2023c). While neither of these currently meet the requirements for being defined as a clear detection, further investigations into the analysis methods, as well as the upcoming combined IPTA data set, are expected to improve on these GWB measurements.

The signal induced by the GWB in pulsar data is expected to be extremely weak, even when correlated over tens of pulsars (e.g. Siemens et al. 2013; Janssen et al. 2015). Therefore, other effects in pulsar timing data must be carefully considered, as they can obscure, or even mimic, a GWB signal (Tiburzi et al. 2016). One of the strongest sources of ‘noise’ in this context are the dispersive delays introduced in the ToAs by the interaction between the radio waves and the ionized medium through which the pulsar signal propagates on its way to the observer (Lentati et al. 2016). These dispersive delays are modelled as having an inverse-squared dependency with the observing frequency, f_{obs} . The dispersive delay on a ToA is expressed as

$$t_{\text{D}} = \frac{1}{K_{\text{D}} f_{\text{obs}}^2} \int_{\ell} n_{\text{e}} d\ell = \frac{1}{K_{\text{D}} f_{\text{obs}}^2} \text{DM}, \quad (1)$$

where $K_{\text{D}} \simeq 2.41 \times 10^{-4} \text{ MHz}^{-2} \text{ pc cm}^{-3} \text{ s}^{-1}$ is a dispersion constant (Manchester & Taylor 1972), and the dispersion measure DM is defined as the free electron number density, n_{e} , integrated over the line of sight to the pulsar, ℓ ; the DM is usually quoted in pc cm^{-3} . As illustrated by equation (1), the dispersion delay is stronger at lower observing frequencies. The noise due to the dispersion delay is dominated by the effects of the turbulent and inhomogeneous ionized interstellar medium (IISM) along the line-of-sight, which can induce fluctuations in the DM of order $10^{-3} \text{ pc cm}^{-3}$ over a time-scale of years (e.g. Keith et al. 2013; Jones et al. 2017; Donner et al. 2020). Several possible mitigating strategies for the turbulent IISM contribution are used throughout the PTAs, such as modelling it as a chromatic red-noise Gaussian process in Fourier space (e.g. Lentati et al. 2014), or using a time-domain piecewise binned model (e.g. the DMX model; NANOGrav Collaboration et al. 2015).

For pulsars with a line-of-sight that passes close to the Sun, the delay induced by the propagation through the solar wind (SW hereafter) is also noticeable in the current quality of pulsar data, inducing DM fluctuations as high as 10^{-3} – $10^{-4} \text{ pc cm}^{-3}$. Furthermore, Tiburzi et al. (2016) showed that, if not carefully considered, the influence of the SW may mimic a GWB signal in PTA-like data, as the SW can create spatial correlations among the pulsars. To account for the influence of the SW in pulsar timing data, PTA collaborations generally use a simple ‘spherical SW’ model, based on the assumption that the number density of ionized electrons varies under a spherically symmetric law away from the Sun, according to the inverse square law (Edwards, Hobbs & Manchester 2006), i.e.

$$n_{\text{e}}(\mathbf{r}) = N_{\text{e}}^{\text{SW}} \left(\frac{1 \text{ au}}{|\mathbf{r}|} \right)^2, \quad (2)$$

where \mathbf{r} is the position vector from the Sun to the point of interest affected by the SW, generally given in astronomical units (au); N_{e}^{SW} is the amplitude of the number density at 1 au, and we generally refer to it as the ‘amplitude of the SW’ in this work. Under the spherical SW assumption, this amplitude N_{e}^{SW} is space invariant (does not depend on the vector \mathbf{r}); thus integrating the number density as per equation (2) over the line-of-sight gives the DM contribution as (e.g.

You et al. 2007)

$$\text{DM}_{\text{sph}}^{\text{SW}} \simeq 4.85 \times 10^{-6} \left(\frac{N_{\text{e}}^{\text{SW}}}{\text{cm}^{-3}} \right) \left(\frac{\pi - \theta}{\sin \theta} \right) \text{pc cm}^{-3}, \quad (3)$$

where θ is the solar elongation angle of the pulsar, i.e. the pulsar-observer-Sun angle, which is minimum at the solar conjunction of the pulsar. Replacing the DM given as per equation (3) into equation (1) gives the time delay of a ToA, due to a spherical SW, which we factorise as

$$t_{\text{D}}^{\text{SW}} = \frac{1}{K_{\text{D}} f_{\text{obs}}^2} N_{\text{e}}^{\text{SW}} \mathcal{S}_{\theta}, \quad (4)$$

where we have summarized some of the physical constants and geometrical dependence in the variable

$$\mathcal{S}_{\theta} \simeq 4.85 \times 10^{-6} \left(\frac{\pi - \theta}{\sin \theta} \right) \text{pc} \simeq \left(\frac{\pi - \theta}{\sin \theta} \right) \text{au}. \quad (5)$$

In the standard approach to mitigating the effect of the SW, the amplitude N_{e}^{SW} in the described spherically symmetric model is constant in time; in e.g. the recent EPTA data set (EPTA Collaboration et al. 2023a), this amplitude is generally kept fixed at 7.9 cm^{-3} , as per Madison et al. (2019). Furthermore, data taken when the pulsar appears $<5^\circ$ away from the Sun in the sky are considered to be poorly described by this simple model and commonly removed (Verbiest et al. 2016).

No model has yet been developed that fully captures the observed impact of the SW, especially on low-frequency pulsar data, as shown by e.g. Tiburzi et al. (2019). You et al. (2007) proposed a model based on the coronal magnetograms derived by the Wilcox Solar Observatory, and the bi-modal nature of the SW, considering the contributions from both a slow (equatorial) and a fast (polar) solar stream (e.g. Coles 1996). You et al. (2007) argued that this two-phase model performed better than the spherical SW for their PPTA observations. However, Tiburzi et al. (2019) found that the spherically symmetric SW model with time-dependent amplitude performed systematically better than the two-phase model in removing the SW contribution in their longer-span and lower-frequency LOFAR data. The apparent discrepancy between these two analyses is thought to be due to either (i) the increased DM precision of the lower observing frequency of the data used in Tiburzi et al. (2019) as compared with that of You et al. (2007); or (ii) a difference in performance of the two-phase model with the heliospheric latitude of the pulsar, since the two papers investigated data from different pulsars.

A clear improvement to current general models is to allow the amplitude of the SW to vary each year (i.e. be time-variable) in the spherically symmetric model. Tiburzi et al. (2021), hereafter referred to as T21, have already shown this to be beneficial and more adequate for low-frequency data taken with the European interferometer LOFAR (Van Haarlem et al. 2013). T21 fit a time-variable spherical SW to each pulsar, and a clear temporal variation was observed in several of the analysed pulsars. Note that previously, Madison et al. (2019) had reported little evidence for long-term variations in the SW density using the NANOGrav 11-yr data set.

In this work, we present a Gaussian-process piecewise approach, implemented as part of the pulsar analysis toolkit `RUN_ENTERPRISE` (Keith, Niřu & Liu 2022), which allows for an automatic time-variable spherical SW fitting in a Bayesian framework, simultaneously with all the other pulsar timing parameters and noise models. Recently, Hazboun et al. (2022) developed several comprehensive Bayesian algorithms that allow fitting for SW across several pulsars simultaneously. They used a uniform-prior piecewise model that globally fits for a spherical SW amplitude in each temporal bin

Table 1. Ecliptic latitudes (elat) and summary of which data sets are available for the three pulsars used in this work; a tick (✓) signifies the data are available, whereas a cross (✗) means they are not. By EPTA availability, we refer specifically to the Data Release 2 data set. The EPTA and InPTA data are combined for PSR J1022+1001.

PSR	elat [°]	EPTA+InPTA	PPTA	LOFAR
J0030+0451	1.45	✓	✗	✓
J1022+1001	−0.06	✓	✓	✓
J2145−0450	5.31	✗	✓	✓

(of e.g. 3 months) of the data, also exploring variations in the exponent of the $1/r^2$ law of equation (2); further, they also explore a model based on globally fitting for continuous Fourier-basis, time-dependent variations in the SW. In this work, we present a different approach to SW Bayesian fitting. We use a simple piecewise algorithm based on a Gaussian process with a Normal-distribution prior to fit for a time-variable amplitude spherical SW in each pulsar individually. The choice of prior allows us to marginalize over the individual yearly amplitudes, reducing the problem to a single additional hyperparameter in the width of the Gaussian prior. This, together with the mathematical implementation that keeps the parameter estimation algebraically linear, makes any additional computational cost negligible.

This article is structured as follows: in Section 2, the main properties of the data sets are summarized in Section 3, we describe the algorithm implemented in our pipeline and simulations; in Section 4, the results of our analysis are discussed; and in Section 5, we summarize our conclusions.

2 DATA SETS

For this work, we selected 3 ms pulsars, which pass in close proximity ($\leq 5.31^\circ$) to the Sun during conjunction, and are included in the study by T21. These are PSRs J0030+0451, J1022+1001, and J2145−0450, and they are present both in LOFAR and PTA observations. Here, we consider data from the recent EPTA+InPTA data combination (EPTA Collaboration et al. 2023a; Tarafdar et al. 2022), as well as from the Second Data Release (DR2) of PPTA (Reardon et al. 2021). Table 1 shows the ecliptic latitudes of the three pulsars of interest, equivalent to the sky-angle between the pulsar and the Sun at conjunction. It also shows which data were available for each pulsar. The data sets, each with different frequency coverage and properties (discussed in more detail in Sections 2.1–2.3), are used separately to construct simulations, and thus assess their sensitivity to the SW. Further, the results of our pipeline are compared across all the data sets, as well as with the independent measurements in T21.

2.1 EPTA (DR2)+InPTA (DR1) combined data set

The Second Data Release (DR2) of the EPTA collaboration is used in this work, augmented with the First Data Release (DR1) of the InPTA (Tarafdar et al. 2022; EPTA Collaboration et al. 2023a). Here, we refer to this combined data set as EPTA+InPTA.

There are five European radio telescopes that provided data for the EPTA DR2 data set, namely: the 100-m Effelsberg Telescope (in Germany), Jodrell Bank Observatory’s 76-m Lovell Telescope (in the United Kingdom), Nançay Radio Observatory’s large Radio Telescope (NRT; in France), the Astronomical Observatory of Cagliari’s 64-m Sardinia Radio Telescope (SRT; in Italy), and the Westerbork Synthesis Radio Telescope (WSRT; in the Netherlands). Moreover,

these telescopes were also used collectively, on a monthly cadence, as the Large European Array for Pulsars (LEAP), equivalent to a 194-m sixth interferometric telescope in the EPTA (Bassa et al. 2016). The EPTA DR2 data set contains observations of 25 ms pulsars, up to 25 yr in length. The large majority of these observations were taken at frequencies in the ‘L-band’ (1–2 GHz) and above, with bandwidths of up to 512 MHz, while there were a limited number of observations centred at lower frequencies of 350 MHz. For a detailed description of the properties of the EPTA DR2 telescopes and observations, see Chen et al. (2021) and EPTA Collaboration et al. (2023a).

The InPTA data set includes observations taken with the upgraded Giant Meterwave Radio Telescope (uGMRT; in India) over a period of 3.5 yr (for a detailed description see Tarafdar et al. 2022). The uGMRT took simultaneous observations at two frequency bands, referred to as ‘B3’ (300–500 MHz), and ‘B5’ (1260–1460 MHz), respectively. These simultaneous observations at different frequencies, as well as the wide bandwidths, make the InPTA observations extremely valuable for measuring the DM influence in pulsar data, including the effect of the SW.

2.2 LOFAR data set

For PSRs J0030+0451 and J2145−0750, the same LOFAR data sets as in T21 are used; while for PSR J1022+1001, the T21 data set was supplemented with the most recent ~ 2 yr of observations for this work. These data were taken with subsets of the International LOFAR telescope (Stappers et al. 2011; Van Haarlem et al. 2013), namely six of the German stations, the Swedish station, and the LOFAR Core in the Netherlands. More than 100 pulsars are regularly observed using this set up. The extremely low observing frequency of the LOFAR instruments, covering a range of roughly 110–190 MHz, makes this telescope incredibly well-suited for studying chromatic effects on pulsar data such as the SW. For more details on the LOFAR data sets, see e.g. Porayko et al. (2019), Donner et al. (2019), Tiburzi et al. (2019).

2.3 PPTA (DR2) data set

We also make use of the open-access PPTA DR2 timing data (Reardon et al. 2021) to validate and compare with the EPTA+InPTA results. This PPTA data set is described in detail in Kerr et al. (2020) and was taken with the Australian 64-m Murriyang Parkes Radio Telescope. It spans ~ 14 yr (2004–2018) of observations of 26 ms pulsars at three frequency bands, roughly centred at 700, 1400, and 3100 MHz. The observing cadence for each pulsar was approximately two–three weeks. The ToAs are available in both sub-banded and band-averaged form, of which we used the latter in this work, for simplicity. In particular, the PPTA DR2 data for PSRs J1022+1001 and J2145−0450 were used to compare SW results with those of the LOFAR and EPTA+InPTA ToAs. The data of PSR J2145−0450 was also supplemented with archival observations of the Parkes telescope since 1994, as published in Manchester et al. (2013); these were, however, much less sensitive to the SW.

3 METHOD

3.1 Modelling

3.1.1 Timing model

In this work, we use open-source pulsar software, namely the Bayesian fitting software ENTERPRISE (Ellis et al. 2019), which

has been integrated into the pulsar analysis toolkit `RUN_ENTERPRISE` (Keith, Nițu & Liu 2022), together with the pulsar timing software `TEMPO2` (Edwards, Hobbs & Manchester 2006).

The Bayesian method allows for simultaneous fitting of deterministic pulsar parameters (such as the spin frequency and its derivatives, position, known binary companions, etc.), which are generally marginalized over, unless of particular interest; and of parameters characterizing stochastic timing noise. In this framework, white noise in timing data is described by fitting the parameters ‘EFAC’ and ‘EQUAD’, which can rescale the ToA error bars, and account for additional white noise, respectively (Edwards, Hobbs & Manchester 2006). Red timing noise is described using Fourier-domain Gaussian-process modelling as described in Lentati et al. (2014). The sinusoidal Fourier-basis components at each Fourier frequency f are multiplied by a set of amplitudes described by a Gaussian process; we refer to the covariance of these amplitudes as the power spectral density. In our fitting, we model the power spectral density prior as a power law, as often used in the pulsar community (Van Haasteren et al. 2011; Lentati et al. 2014). However, for this study, we choose to characterize this power law by an amplitude at a reference frequency of 0.1 yr^{-1} , different from the usual choice of 1 yr^{-1} ; this choice ensures the amplitude is more robust against the yearly periodicity of the solar wind. We therefore characterize the power spectral density model as

$$P(f) = \mathcal{A} \left(\frac{f}{0.1 \text{ yr}^{-1}} \right)^{-\gamma} \text{yr}^3, \quad (6)$$

where \mathcal{A} and γ correspond to the Bayesian hyperparameters characterizing the red timing noise, while the individual Gaussian-process amplitudes are marginalized over. This is used for both achromatic red timing noise, and chromatic DM noise, which has an inverse-square dependence on observing frequency, as per equation (1). For the former, the power amplitude is often written as $\mathcal{A} \equiv A_{\text{red}}^2 / (12\pi^2)$, and $\log_{10}(A_{\text{red}})$ is chosen as the fitting hyperparameter together with the exponent $\gamma \equiv \gamma_{\text{red}}$. For the DM noise, the Fourier components are also proportional to the square-inverse of the observing frequency f_{obs} ; the prior power amplitude is expressed as $\mathcal{A} \equiv A_{\text{DM}}^2$, and $\log_{10}(A_{\text{DM}})$ and $\gamma \equiv \gamma_{\text{DM}}$ are the fitting hyperparameters.

3.1.2 Solar wind fitting using Gaussian processes

We present an implementation of time-variable amplitude spherical SW modelling, which has been incorporated in the `RUN_ENTERPRISE` package and can therefore be included in the simultaneous Bayesian fitting of pulsar timing data, together with the other deterministic pulsar parameters, and the white and red noise parameters.

An approach based on Gaussian processes is used, similar in concept to that used to describe the red noise model as presented in Lentati et al. (2014). The SW signal is expressed, in the time domain, as the sum of independent components, which correspond to the unit-amplitude spherical SW for each solar conjunction, multiplied by a set of amplitudes equivalent to the quantity N_e^{SW} from equation (3). We choose a set of simple piecewise linear (‘triangular’) functions, with centre points at the solar conjunction times of the pulsar, to mathematically represent each yearly SW variation; this follows the idea presented in Keith et al. (2013) for modelling DM noise in general. In practice, the SW contribution to the vector of ToAs is written as

$$\mathbf{t}_{\text{sw}} = \mathbf{n}\mathbf{V}, \quad (7)$$

where \mathbf{n} is a column vector containing the spherical SW amplitudes equivalent to N_e^{SW} for each of the N_c solar conjunctions within the

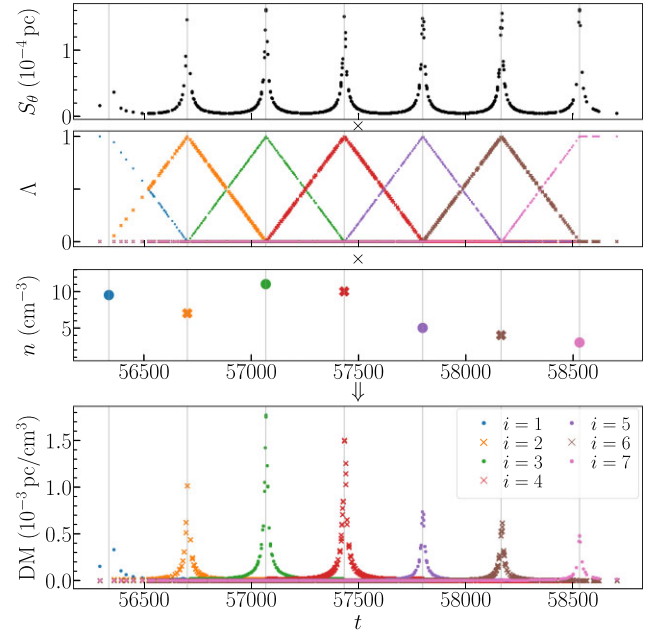


Figure 1. Graphical illustration of the mathematical representation in equations (7)–(9). The quantities in the top three plots (S_θ , Λ , and n) are multiplied to give the DM (bottom plot). Note that we do not show the dependence on observing frequency for simplicity, and $\text{DM} = K_D f_{\text{obs}}^2 \mathbf{n}\mathbf{V}$. Each colour represents a different component, corresponding to each solar conjunction, T_i . The data points represent the ToAs t_j .

data. \mathbf{V} is a matrix of size $(N_c \times N_t)$, where N_t is the number of ToAs, i.e. the size of the column vector \mathbf{t}_{sw} . We define an element of the matrix \mathbf{V} , such that

$$V_{ij} = \frac{1}{K_D f_{\text{obs}}^2} S_{\theta,j} \Lambda \left(\frac{t_j - T_i}{1 \text{ yr}} \right), \quad (8)$$

where $i \in \{1, \dots, N_c\}$ is the matrix row and $j \in \{1, \dots, N_t\}$ the column; $S_{\theta,j}$ is defined as per equation (5) for the solar elongation angle θ_j corresponding to the ToA t_j . T_i is the time of the i th solar conjunction in the data, and Λ is the triangular function, such that by definition

$$\Lambda(q) = \begin{cases} 1 - |q|, & |q| < 1; \\ 0, & \text{otherwise.} \end{cases} \quad (9)$$

Fig. 1 illustrates the steps in creating the SW components in this model.

We assume the distribution of SW amplitudes to be characterized by a Gaussian function, of mean value N_0 , and a standard deviation σ_{sw} . In practice, we therefore express the vector of amplitudes \mathbf{n} as a sum

$$\mathbf{n} = N_0 \mathbf{1} + \delta \mathbf{n}, \quad (10)$$

where N_0 is the mean SW amplitude (i.e. the equivalent of the standard `TEMPO2` parameter ‘NE.SW’), and $\mathbf{1}$ is an N_c -point column vector of ones. $\delta \mathbf{n}$ is therefore the column vector of SW amplitude variations away from the mean. We assume the elements of $\delta \mathbf{n}$ to have mean zero, and their variation with each solar conjunction to be characterized by a Gaussian process, and therefore by a corresponding covariance matrix. In this analysis, we use a simple constant-variance, such that the covariance matrix of the SW amplitude variations $\delta \mathbf{n}$ is a diagonal matrix with all diagonal elements equal to the variance σ_{sw}^2 . The quantity σ_{sw} therefore sets the prior for the amplitudes $\delta \mathbf{n}$, and is a hyperparameter referred to as ‘SW_sigma’ in

the fitting code, and in the output parameter file. Note that one can also choose a different, more complicated covariance informed by physical processes, e.g. one that follows the 11-yr solar cycle. While this may be explored in the future, we believe that the simple approach taken here is suited for capturing the year-by-year variability in amplitude in the current data sets.

Similarly to the approach taken in Fourier-basis red-noise fitting, we marginalize over the Gaussian-process amplitudes \mathbf{n} and only keep σ_{sw} as a Bayesian fitting hyperparameter. As this model keeps the parameter estimation algebraically linear, including the additional time-variable spherical SW model when fitting a pulsar data set is computationally inexpensive, as it only adds one extra hyperparameter to the entire timing model. Furthermore, TEMPO2 can be used after the Bayesian fit performed with the RUN_ENTERPRISE software package to explicitly find the amplitudes of the spherical SW at each solar conjunction, if of interest. This is done using the implemented constraining of the least-squares fitting, which takes into account the fitted σ_{sw} ; for an explanation of how this works, see Appendix A in Keith et al. (2013).

3.2 Simulations

3.2.1 General set up

To test the capabilities of the SW model presented in this work, we created sets of simulated ToAs of both uniform and PTA-like cadence. For all simulations, we started from a set of ‘idealized’ ToAs (characterized by zero residuals) for a chosen pulsar, and added realistic levels of noise, informed by typical values in the observations available. This was done using TEMPO2 plugins, which generate the types of signal discussed in Section 3.1, namely:

- (i) ADDGAUSSIAN to add white noise;
- (ii) ADDREDNOISE to add achromatic red noise, characterized by a power law prior of user-specified amplitude and slope;
- (iii) ADDDMVAR to add DM noise, also characterized by a power-law prior of user-specified amplitude and slope;
- (iv) ADDARBITRARYDM to add the DM influence from a simulated SW (see Section 3.2.2 for more details on simulating realistic SW DM series).

Using the above types of signals, multiple ToA sets were produced, serving various testing purposes, which we describe in Section 4. Multiple Gaussian-process realizations were created for the same characteristic hyperparameters, such that we were also able to check the robustness of our pipeline with repeat measurements.

3.2.2 Simulating the solar wind

We simulate the SW influence on a pulsar as a DM time series that is then added to the total simulated ToAs using the TEMPO2 plugin ADDARBITRARYDM. We consider the variation with respect to the unit-amplitude ($N_e^{\text{SW}} = 1$) spherical SW, i.e. the ratio of the DM series to S_θ , and sample it as a Gaussian process characterized by an exponential-squared kernel (see e.g. Rasmussen & Williams 2006), of the form

$$k(\tau) = A_k \exp\left(-\frac{\tau^2}{2\lambda^2}\right), \quad (11)$$

where τ represents the ‘distance’ between two observing times, A_k is the kernel amplitude, and λ is the metric, i.e. the scale of the correlation within the signal described by this kernel. The exponential-squared kernel is a somewhat arbitrary choice, but this




Figure 2. An example of two realizations of simulated SW using an exponential-squared kernel with amplitude informed by real data. The data points are at uniformly sampled ToA values. The top plot represents the simulated series y , which is then multiplied by the pulsar-specific geometrical factor S_θ to create a DM series, shown in the bottom plot.

simple stationary kernel is widely used and well-suited for describing smooth functions characterized by a single overall metric. In practice, we choose $\lambda = 400$ d; the exact value of this is not strictly relevant, but a value slightly above 1 yr ensures that the simulated signal at each solar conjunction is roughly independent of that at the other solar conjunctions.

The amplitude A_k broadly characterizes the overall variation away from the mean of the signal. To find a realistic value of A_k , we use the SW DM series of PSR J0034–0534 as presented in fig. 3 of T21. We refer the reader to Tiburzi et al. (2019, 2021) for details on how this was obtained. In short, using 5 yr of LOFAR observations, the total DM contribution at each average epoch was estimated. The contribution of the IISM to this DM series was then modelled by a cubic spline with each piece corresponding to a solar conjunction; the SW contribution was simultaneously modelled assuming a spherical SW as in equation (3), while the amplitude was allowed to vary year-by-year. In this work, the DM series after subtracting the cubic spline of the IISM model was used, and we refer to this as $\text{DM}_{\text{Lo}}^{\text{SW}}$. We define a time series $y_0 \equiv \text{DM}_{\text{Lo}}^{\text{SW}}/S_\theta$, such that y_0 only encompasses the estimated variation of the SW with respect to a spherical SW of time-invariant, unit-amplitude. Note that at a time, T_i of the i th solar conjunction in the data, the value of y_0 is equivalent to the SW amplitude n_i . The time series y_0 is then assumed to be described by a smooth function, characterized by an exponential-squared kernel as in equation (11), of metric $\lambda = 400$ d. To find the representative amplitude A_k , we employ the PYTHON library GEORGE to fit this exponential-squared kernel to the y_0 series, and find a maximum-likelihood value of $A_k \approx 3.7 \text{ cm}^{-3}$.

With the fully parametrized exponential-squared kernel, we are able to draw as many new samples of time-series y (of the same type as y_0) as needed, for any pulsar with an available ephemeris. The unit-amplitude spherical SW (i.e. S_θ) of any known pulsar is easily obtainable from the TEMPO2 plugin ‘GENERAL2’ using the pulsar’s ephemeris, such that samples of SW DM series for any known pulsar are simulated as $\text{DM}_{\text{sim}}^{\text{SW}} \equiv y S_\theta$. Fig. 2 shows an example of two simulated SW DM series obtained in this way. Note that because our method allows the spherical SW model to vary in time, the resulting simulated DM series is not strictly symmetric with respect to the solar conjunction peaks. The sampled series y is not a discrete set of amplitudes for each solar conjunction, but rather a continuous

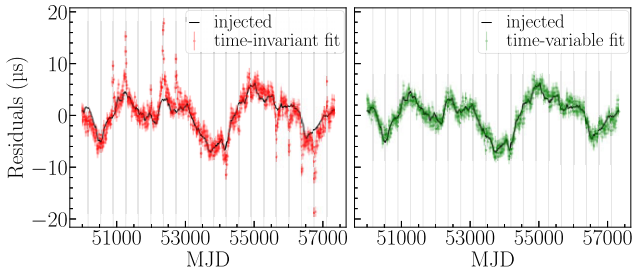


Figure 3. A single realization of simulated residuals for the 300 MHz band, after subtracting the fitted SW variations and achromatic red noise when using the time-invariant (left) and the time-variable (right) SW amplitude model. The solid black line shows the delays at 300 MHz for the injected IISM DM variations. Solar conjunction times are shown by the vertical grey lines.

function that varies with each observation time in a way informed by real data.

4 RESULTS AND DISCUSSION

In this section, we present and discuss the results of our tests of the pipeline and of the sensitivity of available data to SW variations. In all cases, we use the software package `RUN_ENTERPRISE` to simultaneously fit for the standard pulsar deterministic parameters, as well as achromatic and DM red noise with power-law priors. The deterministic parameters are marginalized over, and recovered later if necessary using least-squares fitting in `TEMPO2`; the hyperparameters for the red-noise power-law priors ($\log_{10}(A_{\text{red}})$, γ_{red} , $\log_{10}(A_{\text{DM}})$, γ_{DM}), and the square-root of the variance in SW yearly amplitudes (σ_{sw}) are sampled using a Markov Chain Monte Carlo technique through the `PYTHON` package `EMCEE` (Foreman-Mackey et al. 2013). In this context, we refer to ‘time-invariant’ SW fitting as only fitting for an overall mean amplitude N_0 (i.e. ‘NE_SW’ in `TEMPO2`) of a spherical SW model. This is in contrast to the model described in Section 3.1.2, in which we fit for a ‘time-variable’ spherical SW. In practice, this time-variable spherical SW is described by yearly amplitude variations δn away from an overall mean amplitude N_0 , and we fit for the δn elements, as well as for N_0 .

4.1 Testing the solar wind fitting code

To initially probe the capabilities of our pipeline, we simulate a 20-yr long set of uniform, high-cadence (one observation every 10 days) ToAs of PSR J0034–0534. This testing data set includes simultaneous observations of three common frequency bands, centred at 300, 1440, and 2440 MHz, with an rms of 1 μs on each ToA. While these data are ‘optimistic’ compared with typical pulsar data sets, they provide a useful first test to gain insight into the performance of the SW fitting pipeline. To create the simulated data set, we start from idealized ToAs (of zero residuals), and ‘inject’ white noise variations at the 1- μs level, Gaussian-process achromatic red noise and DM noise with known power-law priors informed by typical real levels of noise seen in the observations, as well as a realistic time-variable SW, using the method described in Section 3.2.2. Multiple realizations of the noise and SW Gaussian-process samples are used for robustness and repeat-measurement checks on the results of the pipeline.

In this section, we aim to explore the proficiency of our pipeline. First, the injected signals are compared to the recovered ones using either the time-invariant or the time-variable SW method. The comparison of the results of the two fitting methods is discussed




Figure 4. The power spectral density (PSD) against Fourier frequency of the recovered DM is shown, averaged over 15 realizations of our simulated data, after fitting for a time-invariant (red) and time-variable (green) SW. The grey dashed line represents the power-law prior of the injected DM Gaussian-process signal.

in Section 4.1.1. Secondly, we use the measurements of 2000 SW amplitudes in order to validate the size of the uncertainties produced by our pipeline; this is described in Section 4.1.2.

4.1.1 Comparison between time-invariant and time-variable fitting

One of the benefits of using simulations is that the quantity of red noise in the pulsar ToAs is known, and we can therefore compare the properties of the recovered signals to those of the injected ones. Here, we run both the time-invariant and the time-variable fitting pipelines on the same set of simulated ToAs, which include both achromatic and DM red noise.

In general, we expect the recovered DM power law to be affected by whether the SW influence is well modelled. Specifically, if only a time-invariant mean amplitude is used to model the SW, any year-by-year variation is likely to be absorbed by the DM noise model instead. This leads to excess power in the measured DM noise at high Fourier frequencies (at 1 yr^{-1} and further harmonics), which in turn causes the power law to appear flatter than the real power-law process. Therefore, unmodelled SW influences may bias the interpretation of the DM spectrum. Furthermore, this could also bias the estimation of other timing parameters, and especially their uncertainties, as the DM power law would not accurately represent the actual noise present in the data.

To investigate the performance of our pipeline compared to the time-invariant fitting in this context, we inspect the low-frequency (300 MHz) residuals after subtracting the maximum-likelihood SW contribution and fitted achromatic noise. Fig. 3 shows these residuals, overlaid with the injected IISM DM variations. When fitting only for a time-invariant SW, we can see additional ‘spikes’ in the residuals around the solar conjunction times, a result of absorbing unmodelled SW into the DM variations, while the time-variable SW fitting is fully consistent with the injected IISM contribution.

The average power spectral densities of these recovered DM influences, as well as the injected power law prior, are shown in Fig. 4. As expected, in the time-invariant SW fit, the absorbed SW influence creates additional power at high frequencies, and therefore flattens the power-law shape compared to the injected prior, while the time-variable fit recovers the injected power-law well. The shape of the average power spectral density in the time-invariant fit shows the effect of the SW 1/yr-frequency and its harmonics, modulated by subtracting a mean-amplitude SW, which creates the observed dips of smaller width.

The injected and recovered red noise hyperparameters for 15 realizations of the simulation are shown in Fig. 5, as defined in the




Figure 5. Measurements obtained from 15 realizations of the uniform simulation as described in Section 4.1.1 are shown. Specifically, these plots show the power-law log-amplitude ($\log A$) and slope (γ) for the DM (top plots) and achromatic red noise (bottom plots). The error bars represent one standard deviation. The vertical dashed lines show the ‘true’ values used in the simulations. In each plot, 15 realizations of the simulation are shown for each of the two fitting modes: the top red circle points represent the measurements obtained from the time-invariant SW fit, while the bottom green squares correspond to the time-variable SW fitting as described in this work.




Figure 6. The injected and recovered SW contributions are compared. The top plot represents the DM series, where the black points show the injected signal, as well as illustrate the cadence of the ToAs, while the green data points are the recovered yearly amplitudes. The bottom plot represents the above DM divided by the geometrical pulsar factor S_θ . The continuous horizontal lines show the mean SW amplitude (N_0) fitted in each of the time-variable (green) and time-invariant (red) model; the shaded region around these lines represent the area within one standard deviation.

power law in equation (6). The achromatic red noise hyperparameters are in agreement with the injected values and equally well recovered in both cases; this may be expected, as the SW influence is an intrinsically chromatic effect. The time-variable SW pipeline also recovers the DM power law well, though this is not the case for the time-invariant SW model. Notably, the time-invariant model leads to a spectral exponent that is significantly flatter, with the recovered slope being of order 10σ smaller than the ‘true’ value. This appears to be a systematic rather than statistical effect, i.e. not dependent on the specific Gaussian-process sample used in the simulation, but rather due to the fitting model.

The recovered SW model from each fitting method can also be compared to the injected SW influence. Fig. 6 shows this for the DM series (top panel), as well as the spherical SW amplitude N_e^{SW} (bottom

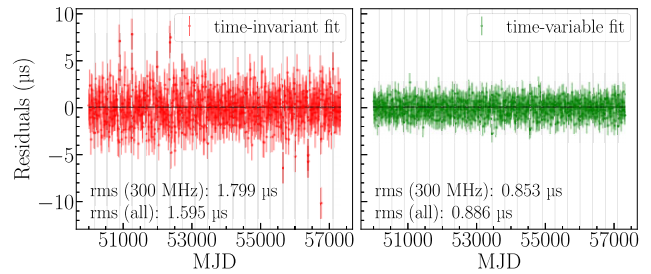


Figure 7. Post-fit residuals after subtracting the entire timing model, including achromatic and DM red noise, are shown for both the time-invariant (left) and time-variable (right) SW amplitude models. Only the data at the 300 MHz observing frequency is shown for clarity, as it shows the largest difference between the two models. The root-mean-square (rms) values for both the 300 MHz and the full data (‘all’) are also presented.

panel). The modelled solar-conjunction amplitudes estimated by our pipeline follow the injected SW well, being within 2σ of all injected values. The time-invariant fit (shown only in the bottom panel of Fig. 6) only models the SW influence with a constant-in-time amplitude, and the uncertainty in the mean amplitude recovered is too small to account for the yearly variation.

We also consider the timing residuals after removing all the modelled signals, including the achromatic red noise and DM noise; these are shown in Fig. 7. Ideally, these fully subtracted residuals should be as close to pure uncorrelated noise as possible. The size of this left-over noise, which can be characterized by its root-mean-square (rms), is also relevant in general when searching for signals in the data, such as, for example, the search for pulsar binary companions, or a GWB signature, for which the lower the noise, the better the chance of detection. Of the two fitting modes used in this test, the yearly variable SW appears to be a better model, as strongly supported by the Bayesian evidences, which yield a natural log-Bayes factor of 160 in favour of the time-variable pipeline; the Bayesian evidences are computed using the nested-sampler DYNesty (Speagle 2020). This effect can also be seen in the fully subtracted residuals (Fig. 7), where the residuals obtained using the time-variable SW fitting qualitatively appear more ‘white’, lacking the unmodelled SW ‘spikes’. However, note that the reduced- χ^2 values of the two whitened time series are indistinguishable (~ 1), likely due to the large number of data points not at solar conjunctions. The level of white noise, quantified here by the rms, is better by $\sim 1 \mu\text{s}$ for the residuals obtained from the time-variable fitting than those from the time-invariant fitting.

4.1.2 Recovered uncertainties

To check the robustness of our results and their uncertainties, 100 sets of ToAs are simulated, all with the same general properties as the data set used in Section 4.1.1. Each of these simulated data sets is created from different Gaussian-process realizations of the same power law prior achromatic red noise and DM noise, as well as different samples of SW variations. The results of running these 100 independent data sets through our pipeline are consistent with the findings in Section 4.1.1 with respect to the injected values.

The 2000 solar conjunctions available in total in the 100 data sets are used to study the statistical properties of the estimated uncertainties on the SW amplitudes. Each measurement of the SW amplitude is normalized with respect to the injected value and the associated uncertainty, such that the equivalent ‘standardized variable’ is computed. The population of these standardized variables

is expected to be normally distributed, with a mean of zero and a standard deviation of one, if the distribution of the measured variable is Gaussian. In the case of a set of uncorrelated measurements, $\{X_i\}$, assumed to be Gaussian distributed, the standardized variable is simply defined as $(X_i - \mu)/\sigma$ for each measurement, where μ and σ are the mean and standard deviation of the Gaussian distribution. In this analysis, however, the measurements of SW amplitudes are correlated through the mean N_0 , such that we use linear algebra to compute the standardized variable values, as follows.

If \mathbf{n} is the column vector of the $m = 2000$ measured SW amplitudes and \mathbf{n}^{inj} is the equivalent column vector of injected SW amplitudes, then we can write

$$\mathbf{n} = \mathbf{n}^{\text{inj}} + \mathbf{e}, \quad (12)$$

where \mathbf{e} is a column vector representing the random variation of \mathbf{n} around \mathbf{n}^{inj} , expected to be normally distributed with mean zero and variance according to the parameter covariance matrix \mathbf{C} of the elements in \mathbf{n} . If \mathbf{C} is known, the vector $\mathbf{e} = \mathbf{n} - \mathbf{n}^{\text{inj}}$ can be ‘whitened’ to disentangle the correlations between different amplitude measurements, and therefore to obtain the column vector of standardized variables \mathbf{z} . This approach is similar to e.g. the whitening presented in Coles et al. (2011) for general pulsar data, and we summarize the steps below. Using the Cholesky decomposition, the parameter covariance matrix is written as $\mathbf{C} = \mathbf{L}\mathbf{L}^T$, where \mathbf{L} is a lower triangular matrix with a real and positive diagonal. The whitening process then yields that the column vector of standardized variables can be computed as

$$\mathbf{z} = \mathbf{L}^{-1}\mathbf{e} = \mathbf{L}^{-1}(\mathbf{n} - \mathbf{n}^{\text{inj}}). \quad (13)$$

Recall that in our pipeline, we fit for the mean SW amplitude N_0 and the deviations from it at each solar conjunction, i.e. $\delta\mathbf{n} = \mathbf{n} - N_0\mathbf{1}$, where in this case, $\mathbf{1}$ is an m -point column vector of ones. The relevant column vector of measured SW parameters is therefore given by the $(m + 1)$ -point column vector $\mathbf{p} = (N_0, \delta\mathbf{n})^T$, with the associated parameter covariance matrix, \mathbf{C}_p . The parameter covariance matrix \mathbf{C}_p is a direct result of the timing analysis, and can be obtained directly from TEMPO2 post-fitting. The whitening covariance matrix \mathbf{C} can then be computed as

$$\mathbf{C} = \mathbf{M}\mathbf{C}_p\mathbf{M}^T, \quad (14)$$

where \mathbf{M} is the transformation matrix from the $(m + 1)$ -point column vector \mathbf{p} to the m -point column vector \mathbf{n} , i.e. $\mathbf{n} = \mathbf{M}\mathbf{p}$, and therefore has the shape

$$\mathbf{M} = \begin{pmatrix} 1 & 1 & 0 & \cdots & 0 \\ 1 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 0 & 0 & \cdots & 1 \end{pmatrix}, \quad (15)$$

where the first column contains exclusively ‘1’s and the rest of the matrix has the shape of an $(m \times m)$ identity matrix. Therefore, the standardized variable \mathbf{z} can be computed as per equation (13) using the Cholesky decomposition of the parameter covariance matrix \mathbf{C} , which can be estimated from the output covariance matrix of the analysis, \mathbf{C}_p , using equation (14) and the shape of the transformation matrix, \mathbf{M} , as given in equation (15).

Fig. 8 shows a histogram of the (normalized) distribution of the measurements of \mathbf{z} for the 2000 solar-conjunction amplitudes fitted with our pipeline. The distribution of this rescaled uncertainty is well described by a Normal distribution, with an Anderson–Darling test equivalent p -value of 0.55 (e.g. D’Agostino 1986). The Normal distribution has mean of $-0.05(2)$ and a standard deviation of $1.09(2)$,




Figure 8. The histogram of the standardized variable (z) obtained from the 2000 recovered solar conjunction amplitude values and uncertainties. The best-fit Gaussian distribution is also shown, with a mean of $-0.05(2)$ and a standard deviation of $1.09(2)$.

not meaningfully distinguishable from a unit-variance zero-mean Normal distribution; we therefore conclude that our uncertainties are well estimated within our assumptions.

4.2 Solar wind fitting in PTA-type simulations

In this section, we aim to establish whether PTA-like data are sensitive to yearly variations in SW away from the mean amplitude. For this purpose, EPTA+InPTA data of PSR J1022+1001 are used as a basis for simulated ToAs, which are then analysed through our pipeline. We use the real cadence and uncertainties of this data set to simulate current observations (Section 4.2.1), as well as the median properties of relevant telescopes to approximate future observations of the EPTA+InPTA (Section 4.2.2).

4.2.1 Current EPTA+InPTA observations

To simulate a set of ToAs with properties characteristic to the EPTA+InPTA data set, we start from the real data of PSR J1022+1001, which shows the highest SW influence of the 25 pulsars included in the EPTA DR2 (as this pulsar has the closest approach to the Sun, having an ecliptic latitude of -0.06°). Idealized ToAs (of zero residuals) are created such that the observing cadence, frequencies, and error bars of the real data are preserved into the simulated ToAs. As before, realistic levels of noise and SW variations, based on observations, are added to these idealised ToAs.

The comparison of the measured red-noise hyperparameters with the injected values is presented in Fig. 9 for both the time-variable and the time-invariant SW fitting. It can be seen that there is no significant difference between the results of the two fitting modes in this case, and that the best-fit values are mostly consistent with the ‘true’ values within their uncertainties. However, the precision of these measurements is, as one would expect, worse than in our previous simulations of uniform cadence and low-frequency discussed in Section 4.1.1.

Fig. 10 shows the fitted SW amplitudes through our pipeline, compared to the injected Gaussian process sample. While our results are generally consistent with the injected values, the uncertainties are large enough that the results for this data set are also consistent with a constant-in-time SW. This outcome is also confirmed through an examination of the residuals after subtracting both the time-variable fitting model and the time-invariant model, which reveal no




Figure 9. Measurements obtained from 15 realizations of the (EPTA+InPTA)-like simulation as described in Section 4.2.1 are shown. See the caption of Fig. 5 for details.




Figure 10. The injected and recovered SW contributions for the EPTA+InPTA type simulations are compared. The black points show the injected signal, as well as the cadence of the ToAs, while the green data points are the recovered yearly amplitudes. The continuous horizontal lines show the mean SW amplitude (N_0) fitted in each of the time-variable (green) and time-invariant (red) models; the shaded region around these lines represents the area within one standard deviation.

significant difference between the properties of the two: the reduced- $\chi^2 \simeq 0.97$ in both cases, while the rms values are 1.37 and 1.33 μs for the time-invariant and time-variable fitting, respectively.

We therefore conclude that the current EPTA+InPTA data set is not sensitive enough to measure yearly changes in the SW amplitude. However, we also note that there is no noticeable disadvantage of including this additional model in the fitting, as it adds no significant time to the computational run, nor does it increase the overall left-over white noise level. Moreover, when there is not sufficient data around a solar conjunction to measure the variation in SW amplitude, the total amplitude value defaults to the mean (as can be seen, especially in the early data in Fig. 10). The lack of sensitivity to measuring SW changes of even the more recent EPTA+InPTA data is indicative of the small number of observations at low frequencies ($\lesssim 1000$ MHz), where the SW influence is larger and therefore easier to quantify, as well as of the relatively small fractional bandwidths of the observations. However, this is expected to change in the future, as more uGMRT (InPTA) data, which includes high-quality simultaneous observations at low frequencies (~ 400 MHz), is combined into the EPTA data set. Note that the current data combination, as used in the simulations in this section, includes some InPTA data, particularly around the solar conjunction around MJD 58360; these InPTA observations are indeed seen to have a plausibly beneficial effect on SW amplitude fitting for that year. Furthermore, the upcoming inclusion of the high-cadence and

Table 2. The properties of the uniform-cadence simulations based on median real-data values.

Observing system	# freq. bands	Freq. range [MHz]	Cadence [d]	rms [μs]
LEAP	1	1440	30	0.25
uGMRT (B3)	32	300–500	12	7
uGMRT (B5)	4	1260–1460	12	22




Figure 11. All plots show the post-fit residuals after removing the fitted achromatic and DM Gaussian-processes red noise. The plots on the left correspond to the results after using the time-invariant SW amplitude fit, while those on the right correspond to the time-variable fit. The top plots show the full simulated data, with the uGMRT-type data in orange and the LEAP-like data in blue. The bottom plots show the same information, but zoomed in and just for the LEAP data. The weighted root-mean-square (W_{rms}) is given in each case for all the data shown in the plot.

very low-frequency LOFAR data into the EPTA data set will unlock next-level sensitivity to the SW.

4.2.2 Simulated ‘future’ EPTA+InPTA observations

With the results of the realistic (EPTA+InPTA)-like simulations in mind (as presented in Section 4.2.1), we aim to investigate the sensitivity of this data set after 10 years of observations at the current observing set up. For this purpose, we create sets of simulations based on the median of recent properties of the LEAP observing system (which is the most sensitive L -band ‘observatory’ in the EPTA data set), and on the median of recent properties of uGMRT observations. The properties of these simulated observations are summarized in Table 2, showing the number of frequency sub-bands for each observing system, the total frequency range, the (uniform) observing cadence, and the (uniform) rms characterizing the system’s white noise. Note that the observations with the two uGMRT frequency bands (B3 and B5) are simultaneous, which is particularly useful for studying chromatic effects. Uniform-cadence observations are used in this case for simplicity; while the sensitivity of any pulsar data to the SW indeed should depend on the density of observations close to its solar conjunction, a detailed study on how this cadence affects measurements and on determining the best observing strategy is beyond the scope of the current work, and is left as a future investigation.

Here, we consider the improvement of the noise budget of the observations when using the time-variable amplitude SW fitting instead of the time-invariant SW fitting. Fig. 11 shows the residuals in both cases, after subtracting all the fitted signals, including the red

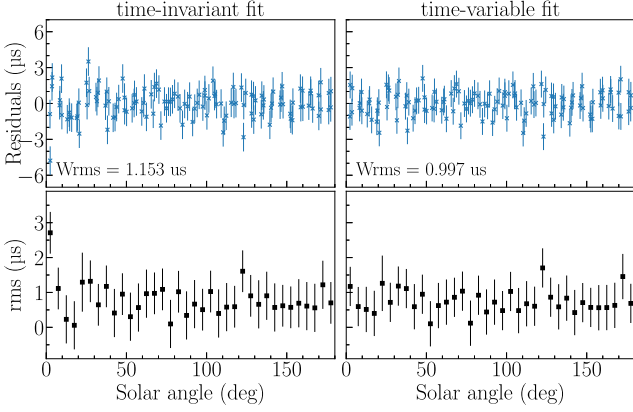


Figure 12. The data shown in the top plots is similar to the LEAP residuals shown in the bottom plots of Fig. 11, but the residuals are plotted against the solar angle, i.e. the sky distance between the pulsar line-of-sight and the Sun. Again, the plots on the left correspond to the time-invariant SW amplitude fit, while the plots on the right show the results of the time-variable fit. The top plots show the residuals for every individual data point, while the bottom plots show the rms of these residuals in bins of 5° solar angle. Of particular interest are the points at low solar angle, i.e. where the pulsar is close to solar conjunction and the SW influence is the largest.

noise processes. The time-variable fit generally performs better in whitening the pulsar ToAs. Qualitatively, there are obvious signatures of the SW that are not removed in the time-invariant case, particularly near the solar conjunction close to MJD 62740, visible both in low-frequency uGMRT data and, to a lesser extent, in L -band LEAP observations. Quantitatively, while the reduced- $\chi^2 \simeq 1$ in both cases as before, the overall left-over noise level is reduced by using our fitting compared to the time-invariant mean: for the full simulated data set, the rms is reduced by nearly $1.3 \mu\text{s}$; while just for the L -band LEAP data, the rms is improved by 160 ns . The plots of Fig. 12 show the same LEAP residuals, but as a function of the solar angle of the pulsar, i.e. its angular distance to the Sun as projected on the sky, for both the time-invariant and time-variable fitting. These plots show that the main difference between the two fitting methods is, as expected, due mainly to those observations closest to the Sun, although even some observations at tens of degrees of solar angle are affected.

Overall, we conclude that in the near future, as more InPTA observations will be combined in the EPTA data sets, the sensitivity to the SW in pulsar data will increase. Moreover, by fitting this data for a time-variable spherical SW, the noise budget is likely to improve, even for L -band observations. This is particularly interesting in the context of e.g. detecting the GWB by PTAs, where a decrease in noise of just a few hundred nanoseconds could be valuable in reaching the target detection significance in light of the recent GWB results (EPTA Collaboration et al. 2023a, b, c).

4.3 Real data

We also test our pipeline on real data, of PSRs J0030+0451, J1022+1001, and J2145–0750, respectively. First, our pipeline’s SW amplitude estimates from independent data sets taken by the observing systems EPTA+InPTA, PPTA, and LOFAR are compared in Section 4.3.1. Secondly, the recovered DM time-series from the fits on LOFAR data are compared to those given in T21, which used the same (while slightly shorter) LOFAR data sets, but a different method; we are therefore able to directly compare SW

Table 3. Table summarizing the uncertainties in the measured SW amplitudes (N_e^{SW}), as well as the mean values (N_0) from the real data sets of the three pulsars discussed, as shown in Figs 13, 14, and 15. The median and minimum values of these measured uncertainties (u.) are quoted in columns 3 and 4, respectively. The values in square brackets represent the median estimated only using ‘newer’ data, where relevant, i.e. after 2004 for PPTA and after 2005 for EPTA. Column 5 summarizes the mean SW amplitude values, with the values in brackets representing one standard deviation.

PSR	Data set	median u. [cm^{-3}]	min u. [cm^{-3}]	N_0 [cm^{-3}]
J0030+0451	EPTA	2.1	1.0	8.8(1.0)
	LOFAR	0.4	0.2	9.5(1.0)
J1022+1001	(E + In)PTA	3.3 [2.4]	1.5	9.9(1.3)
	PPTA	2.0 [1.9]	0.7	11.2(1.0)
	LOFAR	0.3	0.2	9.9(1.0)
J2145–0750	PPTA	2.4 [1.5]	1.0	6.3(1.0)
	InPTA	1.0	0.9	5.0(0.9)
	LOFAR	0.3	0.3	7.1(1.4)




Figure 13. The yearly SW amplitudes for PSR J0030+0451 are shown, resulting from the EPTA (blue dots) and LOFAR (green squares) data, as well as the corresponding values presented in T21 (purple diamonds). The cadences of the observations used are also shown, where ‘E’ stands for the EPTA data and ‘L’ for the LOFAR data. Note that the two x -axes are the same, but presented in both MJD and year for convenience.

results between our pipeline and another, independent method; this is discussed in Section 4.3.2.

4.3.1 Comparison between results from different data sets

Not all telescope data sets were available for each of the three pulsars chosen for this analysis; the data used is summarized in Tables 1 and 3. All data sets are run through the same pipeline, fitting for deterministic parameters, the SW model, white noise, achromatic red noise, and DM noise simultaneously. We compare the yearly SW amplitudes estimated from each separate data set, for each pulsar. This is shown in Figs 13, 14, and 15, while the mean values, as well as the median and minimum of the SW amplitude uncertainties for each data set are also presented in Table 3. The values of SW amplitudes from the independent LOFAR-data analysis from fig. 6 of T21 are also included here for comparison, but are mostly discussed in the following Section 4.3.2.




Figure 14. The yearly SW amplitudes for PSR J1022+1001 are shown, resulting from the EPTA+InPPTA (blue dots), PPTA (orange crosses), and LOFAR (green squares) data, as well as the corresponding values presented in T21 (purple diamonds). The cadences of the observations used are also shown, where ‘E’ stands for the EPTA+InPPTA data, with the cyan lines highlighting the InPPTA data specifically; ‘P’ stands for the PPTA data, and ‘L’ for the LOFAR data. Note that the two x -axes are the same, but presented in both MJD and year for convenience.




Figure 15. The yearly SW amplitudes for PSR J2145–0750 are shown, resulting from the InPPTA (blue dots), PPTA (orange crosses), and LOFAR (green squares) data, as well as the corresponding values presented in T21 (purple diamonds). The cadences of the observations used are also shown, where ‘In’ stands for the InPPTA data, ‘P’ for the PPTA data, and ‘L’ for the LOFAR data. Note that the two x -axes are the same, but presented in both MJD and year for convenience.

First, the results for PSR J0030+0451 are shown in Fig. 13. There is an agreement within uncertainties between the results of the EPTA and LOFAR data sets. As is perhaps expected, the much-lower frequency LOFAR data performs significantly better in estimating the SW amplitudes.

Secondly, Fig. 14 shows the results for PSR J1022+1001. In general, the SW amplitude measurements are consistent between the different observing systems. While the EPTA data set is a few years longer than the others, the early data are sparse enough that only rough estimates of the amplitudes are possible. Broadly, the PPTA data allow for more precise measurements of the SW than the

EPTA data, likely owing to the additional lower-frequency and wider-bandwidth observations in the PPTA data set. Further, the LOFAR data are seen to be much more sensitive to these measurements than the PPTA data. The collection of InPPTA data has started relatively recently, and as such only two amplitude measurements are supplemented by this in the EPTA+InPPTA data combination; these are the last data point and that near MJD 58360. The increased sensitivity of the amplitude measurement which included both EPTA and InPPTA data for the solar conjunction near MJD 58360 hints at the advantage of the additional InPPTA data.

Thirdly, the SW measurements for PSR J2145–0750 are shown in Fig. 15. The early data (before 2004) from the PPTA represent archival data with a low cadence, such that no variation away from the mean-amplitude SW can be measured. The precision of the LOFAR measurements is significantly better than that of the PPTA measurements in the case of PSR J2145–0750 as well. We also note that the recent, roughly 4 years of InPPTA data show an improved precision compared to the PPTA data set for this pulsar, likely due to their high-cadence simultaneous observations.

The observed dip in SW amplitude occurring around the year 2009 (around MJD 55000) for PSR J2145–0750 is coincident with the independently observed low in solar activity (e.g. Hathaway 2015), such that our measurements appear to track the broad behaviour of the solar cycle. The same dip can also be seen, in fact, for one solar conjunction in PSR J0030+0451; however, this is not obviously present in the results of PSR J1022+1001. This suggests that, on one hand, the SW signature in pulsars further away from the ecliptic plane is likely to follow the broad 11-yr solar cycle behaviour. On the other hand, the shape of the variation away from the mean seems harder to predict for pulsars of very low ecliptic latitude (such as PSR J1022+1001). This may be caused, for example, by the slow SW contribution dominating the fast SW around the solar equator, even during the solar cycle maxima, thus showing less of the variation between the two solar cycle stages. The presence of persistent streamers near the solar equator could also induce the kind of DM variations seen in pulsars at very low heliospheric latitudes. Furthermore, the mean SW amplitude of observations of PSR J1022+1001, of $\sim(10 \pm 1) \text{ cm}^{-3}$ is significantly larger than that of PSR J2145–0750, of $\sim(6 \pm 1) \text{ cm}^{-3}$. The very similar range of observing times for these two pulsars suggests that this divergence may be due, as before, to the different ecliptic latitudes, and the corresponding heliospheric latitudes and magnetosphere areas probed. Therefore, we conclude that, in general, a global fit describing the SW properties of multiple pulsars is likely to be advantageous only if the diverse regions of the SW probed by pulsars at varying heliospheric latitudes are carefully considered.

4.3.2 LOFAR results: comparison with Tiburzi et al. (2021)

Finally, the results of our pipeline are compared with the independent analysis in T21, using the same observational LOFAR data available for PSRs J0030+0451, J1022+1001, and J2145–0750. For our results, the data are processed through our pipeline, as with all the previous analyses presented in this work, and the yearly SW amplitudes are estimated. The full DM series are also reconstructed from the fitted parameters, namely as the sum between the SW model, the chromatic red-noise power-law Gaussian process, and the deterministic DM model; the latter is expressed as a polynomial

$$\text{DM}_{\text{det}}(t) = \text{DM}_0 + \text{DM}_1 \times t + \text{DM}_2 \times \frac{t^2}{2}, \quad (16)$$




Figure 16. The recovered DM series and yearly SW amplitudes from our analysis of the LOFAR ToAs (green) are compared with the same quantities as presented by T21 for the same LOFAR data (purple) for PSR J0030+0451. While the top plot shows the broad picture of the DM series, the middle plot is a zoomed-in version of the same data for a clear comparison. For both our results (in green) and the T21 results (in purple), the data points represent the full estimated DM series, which includes the polynomial terms, the Gaussian-process power law (‘PL’), and the SW. The continuous lines represent only the estimated IISM contribution; this was fitted as an average DM_0 plus multiple consecutive cubic terms in T21, and as a quadratic plus a Gaussian process power-law in our analysis. The bottom plot simply illustrates the SW amplitudes, as also shown in Fig. 13, but only for LOFAR data.

where t is a ToA expressed with respect to a chosen epoch (‘DMEPOCH’ in TEMPO2), and the three polynomial coefficients (DM_0 , DM_1 , and DM_2) are equivalent to the ‘DM’, ‘DM1’, and ‘DM2’ fitting parameters in TEMPO2.

Conversely, the analysis in T21 was based on first obtaining a DM value per observation; for more details on how this was performed, we refer the reader to Tiburzi et al. (2019, 2021). To estimate the SW from the DM series, a reference value (equivalent to DM_0) was subsequently subtracted from the DM series. The remaining contributions of the IISM and SW were disentangled and simultaneously modelled using a Bayesian framework for each specific segment of data corresponding to a solar conjunction. The IISM contribution was modelled as a cubic polynomial for each solar-

conjunction segment, and continuity between different segments was insured. In short, the total DM series in the T21 analysis was modelled as a sum between the reference value DM_0 , the IISM cubic-polynomial, and the spherical SW of yearly variable amplitude. Note that the DM series used here for comparison with our full reconstructed DM is the one initially obtained from the observations by T21, prior to the SW analysis.

Figs 16, 17, and 18 show the results of our analyses, as well as those of T21 for comparison, for the three PSRs J0030+0451, J1022+1001, and J2145–0750. Note that more recent LOFAR data were available for PSR J1022+1001 since the T21 work, which are included in our analysis; the additional data that appears here within one solar conjunction either side was discarded in T21 as there were




Figure 17. The same as in Fig. 16, but for PSR J1022+1001.

not enough ToAs to provide a robust estimate by the criteria chosen in that work.

From these plots, we conclude that our estimates of the SW amplitudes are, for the most part, consistent with those of T21, while using an independent Bayesian approach. We note that where there are discrepancies in the SW amplitudes, they appear to be caused by the difference in the estimated IISM contribution. For example, at the solar conjunction near MJD 58000 of PSR J1022+1001 (Fig. 17), our IISM estimate shows an overall flat behaviour, while the T21 analysis presented a higher, cubic-varying DM estimate. This emphasizes an advantage of our pipeline with respect to the T21 method: the IISM contribution to DM is estimated for the entire data set, rather than in segments, which in general ensures a smoother, and likely more realistic behaviour. We also note that, for the specific measurement near MJD 58000 for PSR J1022+1001, the observed difference in IISM estimates of roughly $10^{-4} \text{pc cm}^{-3}$ would only correspond to a direct difference in the SW amplitude (N_e^{SW}) of roughly 10^{-2}cm^{-3} ; the actual measured difference is, however, of

order $\sim 1 \text{cm}^{-3}$. Nevertheless, we believe that, since this pulsar has a very small ecliptic latitude, the tails of the SW influence may impact the shape and peak of the SW model to a larger extent than would be trivially expected.

Lastly, it is important to acknowledge that the spherical SW model, even while allowing its amplitude to change each year, is not sufficient to fully describe the observed SW influence on these data, as also observed previously in, e.g. Tiburzi et al. (2016). This is seen in our results, as the IISM contribution to the DM clearly absorbed some asymmetric features close to the solar conjunction, which are highly likely due to the solar influence. Therefore, any analysis assuming a spherically symmetric SW, which includes that presented in this work, would not allow for an entirely robust studying of the SW influence in pulsar data. The pipeline developed and presented here, however, can create a basis that can be straight-forwardly adapted in future work to improve on this and include various additional models, such as e.g. different piecewise components for before and after the solar conjunctions, allowing the tails of the SW variation (i.e. the




Figure 18. The same as in Fig. 16, but for PSR J2145–0750.

pulsar ingress and egress from the Sun) to change independently of each other.

5 CONCLUSIONS

We have implemented a computationally inexpensive, linear Gaussian-process piecewise Bayesian approach to fit pulsar ToAs for a spherical SW of yearly time-variable amplitude, simultaneously to all the other pulsar timing and noise parameters; this is available through the pulsar analysis toolkit `RUN_ENTERPRISE`. We have explored the functionality of this fitting pipeline using simulations and found that it performs as expected, and better than the currently widely used time-invariant spherical SW, particularly in data of low-frequency and large fractional bandwidths.

Using simulations, we found that the current EPTA+InPTA data set is not yet sensitive enough to measure variations of the spherical SW amplitude. However, assuming that the current observing strategy of uGMRT will continue, future EPTA+InPTA data will

have increased sensitivity to the SW, such that using our time-variable fitting could improve the rms (white-noise levels) of the residuals by at least ~ 200 ns at L -band, which may help with the noise budget for extremely sensitive experiments such as the search for the GWB. Indeed, recovering the SW influence in EPTA+InPTA pulsar data will depend on the cadence of observations near the solar conjunctions, but a study of these effects is beyond the scope of this work, and is left as a future investigation.

We also applied our pipeline to real data of three pulsars, i.e. J0030+0451, J1022+1001, and J2145–0450, that are known to show the influence of the SW in their ToAs, and were also part of the previous study by T21. The SW amplitudes found from individual fitting of pulsar data from the EPTA+InPTA and/or the PPTA were compared with the results from LOFAR data through our pipeline, as well as the independent results from T21. This showed that SW amplitudes found from different data sets were mostly consistent with each other and that, as expected, where lower-frequency and wider-bandwidth data were present, the uncertainties were reduced.

Furthermore, the variation in the fitted SW amplitudes for the pulsar of the highest ecliptic latitude in this study (J2145–0450; $\text{elat} = 5.31^\circ$) roughly followed the 11-yr solar cycle, while for the pulsar of very low ecliptic latitude (J1022+1001; $\text{elat} = -0.06^\circ$), the amplitude variation generally did not seem to correlate with the long-term solar cycle. This hints that a global SW fit may be more beneficial if, e.g. it is performed in slices of ecliptic latitude. A larger study would clarify if this is indeed a wide systematic effect.

The DM series, including the SW effect, were also compared for the same LOFAR data as recovered from our pipeline and as used in T21. We found that where there was a difference between our SW amplitude and that in T21, this was likely due to the estimated IISM background; we believe our simultaneous fitting of deterministic and noise pulsar components produces a more plausible shape for this background than the consecutive cubic fits as used in T21. However, we conclude that even with the addition of the time-variable amplitude, the spherical SW model is not enough to fully account for the SW influence in LOFAR data. The left-over SW absorbed in the background DM series suggests that perhaps using an asymmetric model, with different components pre- and post-solar conjunction may be beneficial. The piecewise framework presented in this work could be modified for this, or for an alternate way to further study SW models.

ACKNOWLEDGEMENTS

ICN was supported by the STFC doctoral training grant ST/T506291/1. The authors thank the InPTA collaboration for providing the InPTA DR1 data set.

Part of the EPTA data used in this work is based on observations with the 100-m telescope of the Max-Planck-Institut für Radioastronomie (MPIfR) at Effelsberg in Germany. Pulsar research at the Jodrell Bank Centre for Astrophysics and the observations using the Lovell Telescope are supported by a Consolidated Grant (ST/T000414/1) from the UK's Science and Technology Facilities Council (STFC). The Nançay radio Observatory is operated by the Paris Observatory, associated with the French Centre National de la Recherche Scientifique (CNRS), and partially supported by the Region Centre in France. We acknowledge financial support from 'Programme National de Cosmologie and Galaxies' (PNCG), and 'Programme National Hautes Energies' (PNHE), funded by CNRS/INSU-IN2P3-INP, CEA, and CNES, France. We acknowledge financial support from Agence Nationale de la Recherche (ANR-18-CE31-0015), France. The Westerbork Synthesis Radio Telescope is operated by the Netherlands Institute for Radio Astronomy (ASTRON) with support from the Netherlands Foundation for Scientific Research (NWO). The Sardinia Radio Telescope (SRT) is funded by the Department of University and Research (MIUR), the Italian Space Agency (ASI), and the Autonomous Region of Sardinia (RAS) and is operated as a National Facility by the National Institute for Astrophysics (INAF).

The LOFAR data used in this paper is based on data obtained with: i) the German stations of the International LOFAR Telescope (ILT), constructed by ASTRON (Van Haarlem et al. 2013) and operated by the German Long Wavelength (GLOW) consortium (glowconsortium.de) during station-owners time and proposals LC0.014, LC1.048, LC2.011, LC3.029, LC4.025, LT5.001, LC9.039, LT10.014; LC12.019; LC13.007; LT14.006 ii) the LOFAR core, during proposals LC0.011, DDT0003, LC1.027, LC1.042, LC2.010, LT3.001, LC4.004, LT5.003, LC9.041, LT10.004, LPR12.010; iii) the Swedish station of the ILT during observing proposals carried out from May 2015 to January 2018.

We made use of data from the Effelsberg (DE601) LOFAR station funded by the Max-Planck-Gesellschaft; the Unterweilenbach (DE602) LOFAR station funded by the Max-Planck-Institut für Astrophysik, Garching; the Tautenburg (DE603) LOFAR station funded by the State of Thuringia, supported by the European Union (EFRE) and the Federal Ministry of Education and Research (BMBF) Verbundforschung project D-LOFAR I (grant 05A08ST1); the Potsdam (DE604) LOFAR station funded by the Leibniz-Institut für Astrophysik Potsdam; the Jülich (DE605) LOFAR station supported by the BMBF Verbundforschung project D-LOFAR I (grant 05A08LJ1); and the Norderstedt (DE609) LOFAR station funded by the BMBF Verbundforschung project D-LOFAR II (grant 05A11LJ1). The observations of the German LOFAR stations were carried out in the stand-alone GLOW mode, which is technically operated and supported by the Max-Planck-Institut für Radioastronomie, the Forschungszentrum Jülich and Bielefeld University. We acknowledge support and operation of the GLOW network, computing, and storage facilities by the FZ-Jülich, the MPIfR, and Bielefeld University and financial support from BMBF D-LOFAR III (grant 05A14PBA), D-LOFAR IV (grants 05A17PBA and 05A17PC1), and D-LOFAR 2.0 (grant 05A20PB1), and by the states of Nordrhein-Westfalia and Hamburg. We acknowledge support from Onsala Space Observatory for the provisioning of its facilities/observational support. The Onsala Space Observatory national research infrastructure is funded through Swedish Research Council grant No 2017–00648.

DATA AVAILABILITY

The EPTA DR2 data underlying the work in this paper are available at <https://doi.org/10.5281/zenodo.8164424>, as per EPTA Collaboration et al. (2023a). The InPTA DR1 data are available at <https://github.com/inpta/InPTA.DR1>, as per Tarafdar et al. (2022). The PPTA DR2 data are available at <https://doi.org/10.25919/cx59-a798>, as per Reardon et al. (2021). The data from the LOFAR core can be downloaded from the Long Term Archive (LTA; lta.lofar.eu). The LOFAR data taken with the GLOW stations are accessible upon demand by emailing any of the first three authors.

REFERENCES

- Agazie G. et al., 2023, *ApJ*, 951, L8
 Backer D. C., Kulkarni S. R., Heiles C., Davis M. M., Goss W. M., 1982, *Nature*, 300, 615
 Bassa C. G. et al., 2016, *MNRAS*, 456, 2196
 Burke-Spolaor S. et al., 2019, *A&A Rev.*, 27, 5
 Chen S. et al., 2021, *MNRAS*, 508, 4970
 Coles W. A., 1996, *Ap&SS*, 243, 87
 Coles W., Hobbs G., Champion D. J., Manchester R. N., Verbiest J. P. W., 2011, *MNRAS*, 418, 561
 D'Agostino R. B., 1986, *Goodness-of-Fit Techniques*. Marcel Dekker, New York
 Demorest P. B. et al., 2013, *ApJ*, 762, 94
 Desvignes G. et al., 2016, *MNRAS*, 458, 3341
 Donner J. Y. et al., 2019, *A&A*, 624, A22
 Donner J. Y. et al., 2020, *A&A*, 644, A153
 EPTA Collaboration et al., 2023a, *A&A*, 678, A48
 EPTA Collaboration et al., 2023b, *A&A*, 678, A49
 EPTA Collaboration et al., 2023c, *A&A*, 678, A50
 Edwards R. T., Hobbs G. B., Manchester R. N., 2006, *MNRAS*, 372, 1549
 Ellis J. A., Vallisneri M., Taylor S. R., Baker P. T., 2019, ENTERPRISE: Enhanced Numerical Toolbox Enabling a Robust Pulsar Inference Suite. record (ascl:1912.015)

- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *PASP*, 125, 306
- Hathaway D. H., 2015, *Living Rev. Sol. Phys.*, 12, 4
- Hazboun J. S. et al., 2022, *ApJ*, 929, 39
- Janssen G. et al., 2015, in *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*. p. 37
- Jones M. L. et al., 2017, *ApJ*, 841, 125
- Joshi B. C. et al., 2022, *IA&A*, 43, 98
- Keith M. J. et al., 2013, *MNRAS*, 429, 2161
- Keith M. J., Nişu I. C., Liu Y., 2022, *run_enterprise*
- Kerr M. et al., 2020, *PASA*, 37, e020
- Lentati L., Alexander P., Hobson M. P., Feroz F., van Haasteren R., Lee K. J., Shannon R. M., 2014, *MNRAS*, 437, 3004
- Lentati L. et al., 2016, *MNRAS*, 458, 2161
- Madison D. R. et al., 2019, *ApJ*, 872, 150
- Manchester R. N., 2017, *J. Phys. Conf. Ser.*, 932, 012002
- Manchester R. N., Taylor J. H., 1972, *Astrophys. Lett.*, 10, 67
- Manchester R. N. et al., 2013, *PASA*, 30, e017
- NANOGrav Collaboration et al., 2015, *ApJ*, 813, 65
- Porayko N. K. et al., 2019, *MNRAS*, 483, 4100
- Rasmussen C. E., Williams C. K. I., 2006, *Gaussian Processes for Machine Learning*. The MIT Press, Cambridge, MA
- Reardon D. J. et al., 2021, *MNRAS*, 507, 2137
- Reardon D. J. et al., 2023, *ApJ*, 951, L6
- Siemens X., Ellis J., Jenet F., Romano J. D., 2013, *Class. Quant. Gravity*, 30, 224015
- Speagle J. S., 2020, *MNRAS*, 493, 3132
- Stappers B. W. et al., 2011, *A&A*, 530, A80
- Tarafdar P. et al., 2022, *PASA*, 39, e053
- Tiburzi C., 2018, *PASA*, 35, e013
- Tiburzi C. et al., 2016, *MNRAS*, 455, 4339
- Tiburzi C. et al., 2019, *MNRAS*, 487, 394
- Tiburzi C. et al., 2021, *A&A*, 647, A84
- Van Haarlem M. P. et al., 2013, *A&A*, 556, A2
- Van Haasteren R. et al., 2011, *MNRAS*, 414, 3117
- Verbiest J. P. W. et al., 2016, *MNRAS*, 458, 1267
- You X. P., Hobbs G. B., Coles W. A., Manchester R. N., Han J. L., 2007, *ApJ*, 671, 907
- ¹Jodrell Bank Centre for Astrophysics, Department of Physics and Astronomy, The University of Manchester, Manchester M13 9PL, UK
- ²INAF — Osservatorio Astronomico di Cagliari, via della Scienza 5, I-09047 Selargius (CA), Italy
- ³University of Hamburg, Gojenbergsweg 112, D-21029 Hamburg, Germany
- ⁴Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, D-53121 Bonn, Germany
- ⁵Kavli Institute for Astronomy and Astrophysics, Peking University, Beijing 100871, China
- ⁶Laboratoire de Physique et Chimie de l'Environnement et de l'Espace LPC2E UMR7328, Université d'Orléans, CNRS, F-45071 Orléans, France
- ⁷Observatoire Radioastronomique de Nançay, Observatoire de Paris, Université PSL, CNRS, Université d'Orléans, F-18330 Nançay, France
- ⁸Ruhr University Bochum, Faculty of Physics and Astronomy, Astronomical Institute (AIRUB), D-44780 Bochum, Germany
- ⁹Thüringer Landessternwarte, Sternwarte 5, D-07778, Tautenburg, Germany
- ¹⁰ASTRON, Netherlands Institute for Radio Astronomy, Oude Hoogeveensedijk 4, NL-7991 PD, Dwingeloo, the Netherlands
- ¹¹Department of Astrophysics/IMAPP, Radboud University Nijmegen, P.O. Box 9010, NL-6500 GL Nijmegen, the Netherlands
- ¹²School of Physics, Trinity College Dublin, College Green, Dublin 2, D02 PN40, Ireland
- ¹³Fakultät für Physik, Universität Bielefeld, Universitätsstr. 25, D-33615, Bielefeld, Germany
- ¹⁴Shanghai Astronomical Observatory, Chinese Academy of Sciences, 80 Nandan Road, Shanghai 200030, China
- ¹⁵E. A. Milne Centre for Astrophysics, University of Hull, Cottingham Road, Kingston-upon-Hull, HU6 7RX, UK
- ¹⁶Centre of Excellence for Data Science, AI and Modelling (DAIM), University of Hull, Cottingham Road, Kingston-upon-Hull, HU6 7RX, UK
- ¹⁷Dipartimento di Fisica 'G. Occhialini', Università degli Studi di Milano-Bicocca, Piazza della Scienza 3, I-20126 Milano, Italy
- ¹⁸INFN, Sezione di Milano-Bicocca, Piazza della Scienza 3, I-20126 Milano, Italy
- ¹⁹Laboratoire Univers et Théories, Observatoire de Paris, Université PSL, Université de Paris Cité, CNRS, F-92190 Meudon, France
- ²⁰Leibniz-Institut für Astrophysik Potsdam, D-14482 Potsdam, Germany

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.