

Equation-of-state-informed pulse profile modeling

Mariska Hoogkamer¹, Nathan Rutherford², Daniela Huppenkothen¹, Benjamin Ricketts^{1,3},
 Anna L. Watts¹, Melissa Mendes^{4,5,6}, Isak Svensson^{4,5,6}, Achim Schwenk^{4,5,6},
 Michael Kramer⁷, Kai Hebeler^{4,5,6}, Tuomo Salmi⁸, and Devarshi Choudhury¹

¹*Anton Pannekoek Institute for Astronomy, University of Amsterdam,
 Science Park 904, 1098XH Amsterdam, The Netherlands*

²*Department of Physics and Astronomy, University of New Hampshire,
 Durham, New Hampshire 03824, USA*

³*SRON, Niels Bohrweg 4, 2333 CA Leiden, The Netherlands*

⁴*Technische Universität Darmstadt, Department of Physics, 64289 Darmstadt, Germany*

⁵*ExtreMe Matter Institute EMMI, GSI Helmholtzzentrum für Schwerionenforschung GmbH,
 64291 Darmstadt, Germany*

⁶*Max-Planck-Institut für Kernphysik, Saupfercheckweg 1, 69117 Heidelberg, Germany*

⁷*Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, D-53121 Bonn, Germany*

⁸*Department of Physics, University of Helsinki, P.O. Box 64, FI-00014 University of Helsinki, Finland*



(Received 31 October 2025; accepted 2 December 2025; published 24 March 2026; corrected 1 April 2026)

NICER has enabled mass-radius inferences for pulsars using pulse profile modeling (PPM), providing constraints on the equation of state (EOS) of cold, dense matter. To date, PPM and EOS inference have been carried out as two separate steps, with the former using EOS-agnostic priors. This approach has several drawbacks. Ideally, one would perform a fully hierarchical Bayesian inference where the pulse profile and EOS model parameters are jointly fit, but implementing such a framework is complex and computationally demanding. Here, we present an intermediate solution introducing an EOS-informed prior on mass radius into the existing PPM pipeline using normalizing flows. By focusing on the parameter space consistent with certain EOSs, this approach both tightens constraints on neutron star parameters while reducing computational costs and requiring minimal additional implementation effort. We test this approach on two pulsars, PSR J0740 + 6620 and PSR J0437-4715, and with two EOS model families: a model based on the speed of sound inside the neutron star interior (CS) and a piecewise-polytropic (PP) model. Both EOS models implement constraints from chiral effective field theory calculations of dense matter. For both pulsar datasets, the inferred radius credible intervals are narrower than in the EOS-agnostic case, with CS favoring smaller radii and PP favoring larger radii. For PSR J0437-4715, the EOS-informed priors reveal a new, more extreme geometric mode that is statistically favored but physically questionable. Including the PPM posteriors in the subsequent EOS inference further tightens the mass-radius posteriors through the chiral effective field theory constraints. However, there is also a sensitivity to the high-density extensions, where the PP (CS) model produces a shift toward larger (smaller) radii and corresponding stiffening (softening) of the pressure-energy density relation.

DOI: [10.1103/z2hd-w9sy](https://doi.org/10.1103/z2hd-w9sy)

I. INTRODUCTION

Neutron stars, with their extreme compactness, provide a unique opportunity to study the behavior of dense nuclear matter and constrain the equation of state (EOS) (see, e.g., [1–4]). Observations of x-ray emission from their surfaces serve as a powerful tool for probing their internal

structure, thereby placing constraints on the EOS of cold, dense matter.

The Neutron Star Interior Composition Explorer (NICER; [5]), onboard the International Space Station, has played a pivotal role in this effort by detecting soft thermal x-rays from a specific class of neutron stars: rotation-powered millisecond pulsars. These pulsars exhibit rotationally modulated, pulsed x-ray emission believed to originate from hot spots at the magnetic poles where return currents deposit heat (see, e.g., [6–8]). The x-ray signal is analyzed using the pulse profile modeling (PPM) technique (e.g., [9–11] and references therein), which incorporates relativistic effects arising from the neutron star's rapid spin

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

and intense gravitational field. This approach enables precise measurements of neutron star masses and radii.

To date, NICER has enabled mass and radius inferences for five pulsars: PSR J0030 + 0451 [12–15], PSR J0437-4715 [16], PSR J0740 + 6620 [14,17–22], PSR J1231-1411 [23,24], and PSR J0614-3329 [25]. The resulting mass-radius (M - R) constraints have been widely used to constrain the EOS of cold, dense matter (see, e.g., [26–30]).

The X-ray Pulse Simulation and Inference (X-PSI; [31]) code is an open-source software package for PPM and Bayesian statistical inference.¹ In M - R inference to date, X-PSI has adopted an EOS-agnostic approach, following the methodology in, e.g., [32,33], selecting joint flat priors on mass and radius, and modifying the mass prior when informative radio-timing measurements are available. EOS inference, using the open-source software NEoST [34],² is then performed as a separate analysis step. The main motivation for this two-step approach is to provide the community with posterior samples in the M - R plane that can be studied using any EOS model afterward. However, this strategy has some practical drawbacks:

- (1) High computational costs: The accuracy and precision of the PPM-inferred neutron star properties depend on factors such as data quality and the availability of prior information. When these are limited, the resulting parameter space becomes broad or difficult to explore efficiently. For instance, inference for PSR J0437-4715 required ~ 783 k core hours to achieve convergence [16].
- (2) Sampling of implausible parameter space: The use of flat priors permits sampling of regions, such as very compact solutions, that are already ruled out by physical EOS constraints, yet these regions incur very expensive multiple-image ray-tracing computations.
- (3) Difficulty with multimodality: A source like PSR J0030 + 0451 [14,15] exhibits complex, multimodal solutions that are difficult to sample thoroughly within reasonable compute time with flat priors. Imposing realistic priors that exclude regions already excluded by EOS constraints will focus sampling only on physically meaningful regions of parameter space, and potentially break multimodality and reduce sampling time.
- (4) Because M and R are generally correlated with geometry parameters (e.g., spot size and location), imposing more physically meaningful priors will also lead to inferred geometries that are more realistic and better constrained, in contrast to the current implementation.

In an ideal scenario, one would perform fully hierarchical Bayesian inference where the pulse profile and EOS

parameters are jointly fit. While this would solve the aforementioned issues, the implementation of such a framework is very complex, time consuming, and running it could be computationally intractable.

At present, the X-PSI and NEoST packages are two distinct frameworks, each focusing on different aspects of the full Bayesian model, namely PPM and EOS inference, respectively. PPM is carried out for individual stars inferring parameters such as mass and radius. EOS parameters are introduced only in the second stage, along with stellar central densities, when information from multiple stars is combined. Properly unifying the PPM and EOS inference would substantially increase the number of parameters, as it would require simultaneous inference of neutron star properties, EOS parameters, and central density even for a single star. This framework would eventually make it possible to carry out PPM for multiple stars at the same time. On the other hand, such a joint inference could benefit from Bayesian shrinkage: The requirement that all stars share the same EOS may help break certain parameter degeneracies, potentially making the overall sampling process more efficient.

As an intermediate step, in this work, we introduce an EOS-informed M - R prior into the existing PPM pipeline, while leaving the subsequent EOS inference unchanged. This approach requires minimal additional implementation effort but still constrains the PPM step to regions of parameter space consistent with prior EOS information, thereby providing a way to obtain tighter constraints on neutron star parameters through more physically informed priors. In addition, it helps reduce computational costs by limiting the parameter space that must be sampled, which in turn might lead to quicker convergence. Finally, leaving the PPM and EOS inference as two separate steps has the additional advantage of avoiding having to deal with a substantial increase in the number of parameters as would be the case when unifying the two.

To inform the PPM pipeline of the underlying cold, dense matter EOS prior, we use the two EOS model families considered in Rutherford *et al.* [26]. At nuclear densities beyond the Baym-Pethick-Sutherland (BPS) crust [35], these are based on next-to-next-to-next-to-leading-order (N^3 LO) chiral effective theory calculations of matter in beta equilibrium from Keller *et al.* [36] up to $1.5n_0$, where $n_0 = 0.16 \text{ fm}^{-3}$ is the nuclear saturation density. The EOS prior thus builds in knowledge of the crust and of the EOS at nuclear densities based on our modern understanding of nuclear forces. At higher densities, we use two possible model extensions based on the speed of sound inside the neutron star interior (CS) [37] or based on a piecewise-polytropic (PP) model [38]. In this work, we explore the impact of both prior EOS ensembles on the PPM-derived M - R and geometry posteriors of PSR J0740 + 6620 and PSR J0437-4715, and the effects of the derived M - R posteriors on the subsequent EOS inference analysis.

¹<https://github.com/xpsi-group/xpsi>.

²<https://github.com/xpsi-group/neost>.

The remainder of this paper is organized as follows. In Sec. II A, we present the ideal case of a full hierarchical Bayesian model for the simultaneous inference of neutron star and EOS properties. Section II B introduces our intermediate approach that incorporates EOS prior information to constrain neutron star properties, followed by the EOS inference framework using EOS-prior-informed PPM M – R posteriors. Sections II C and II D describe the x-PSI PPM analysis pipeline and the use of normalizing flows to model the M – R prior derived from the two EOS families. In Sec. III, we summarize the datasets for PSR J0740 + 6620 and PSR J0437–4715, and outline the prior intervals for the PP and CS EOS model parameters used in this analysis. Section IV presents the EOS-prior-informed PPM posteriors and the resulting constraints on the cold, dense matter EOS. Finally, in Sec. V, we discuss our conclusions and directions for future work.

II. METHODOLOGY

Section II A outlines the ideal scenario of a full hierarchical Bayesian model for the joint inference of neutron star and EOS parameters. In Sec. II B, we introduce our intermediate method—that does not require substantial implementation effort—which incorporates EOS prior information to constrain neutron star properties and enables EOS inference based on EOS-prior-informed PPM M – R posteriors. Sections II C and II D detail the x-PSI PPM analysis pipeline and the implementation of normalizing flows to represent the M – R prior derived from the two EOS models, respectively. Complete information of each PPM and EOS inference run, including the exact x-PSI and NEoS versions, data products, posterior sample files, and all the analysis files are available in a Zenodo repository [39].³

A. Full hierarchical Bayesian model

Currently, posterior densities are derived for the parameters of each neutron star independently, and later, the M – R posteriors are used as “data” in EOS inference. However, under the assumption that all neutron stars adhere to the same EOS, it is much more informative to infer both global EOS parameters and the local parameters for all available neutron stars simultaneously, using a hierarchical Bayesian model. A visual representation of this full model, using probabilistic graphical models, is shown in Fig. 1. A key advantage of this approach is the occurrence of Bayesian shrinkage: By constraining all neutron stars to lie on a common EOS curve, we should obtain tighter constraints on the M – R relation and the EOS parameters.

We begin by representing the pulse profile data as $\{D_{k,j}\}_{k,j=1}^{K,J}$, where $D_{k,j}$ denotes the observed data for the




FIG. 1. Probabilistic graphical model illustrating the hierarchical Bayesian framework for the joint inference of neutron star properties and EOS parameters. Arrows denote statistical dependencies between variables. The shaded gray circle represents observed variables, while double circles indicate deterministic variables. The box represents a summation that applies to all parameters contained within it. See Table I for a detailed description of all variables.

k th neutron star with K the total number of neutron stars included in the analysis, and the j th data point with J the total number of data points in each pulse profile. See Table I for an overview of all parameter definitions.

Next, we specify the prior distribution for the parameters not directly affected by the EOS, such as the inclination angle and the hot spot geometry parameters:

$$p(\{\theta_k\}_{k=1}^K) = \prod_{k=1}^K p(\theta_k), \quad (1)$$

with the parameter θ_k not affected by the EOS, assuming independence between different neutron star sources.

In contrast, the neutron star parameters that are affected by the EOS are governed by a deterministic relationship. Using $(\varphi, \varepsilon_{c,k})$, the vector containing the EOS parameters and the central density of the k th neutron star, respectively, and solving the Tolman-Oppenheimer-Volkoff (TOV; see [40,41]) equations, the radius R_k is determined by the M – R relationship of the EOS for a given sample from the prior for M_k . M_k is sampled from a mass prior based on radio-timing observations (when available), or from a uniform prior (when such observations are unavailable), and R_k is

³<https://doi.org/10.5281/zenodo.17257402>.

TABLE I. Definitions for each parameter of the statistical model.

Parameter	Description
$\{D_{k,j}\}_{k,j=1}^{K,J}$	Observed pulse profile data for neutron stars $k = 1, \dots, K$ for each data point $j = 1, \dots, J$
$\{\theta_k\}_{k=1}^K$	Parameters not affected by the EOS
$\{M_k\}_{k=1}^K$	Gravitational mass for each neutron star k
$\{R_k\}_{k=1}^K$	Equatorial radius for each neutron star k
φ	EOS parameters
$\{\varepsilon_{c,k}\}_{k=1}^K$	Central energy density for each neutron star k

determined using φ , $\varepsilon_{c,k}$, and M_k . In this setup, the mass M_k is treated as an independent variable, and we can express its prior as

$$p(\{M_k\}_{k=1}^K) = \prod_{k=1}^K p(M_k), \quad (2)$$

assuming independence between different neutron stars.

The prior on the EOS parameters φ is given by

$$p(\varphi). \quad (3)$$

Similarly, the prior on the central energy density $\varepsilon_{c,k}$ is given by

$$p(\{\varepsilon_{c,k}\}_{k=1}^K) = \prod_{k=1}^K p(\varepsilon_{c,k}), \quad (4)$$

again assuming independence across neutron stars. We note that, strictly speaking, $\{\theta_k\}_{k=1}^K$, $\{\varepsilon_{c,k}\}_{k=1}^K$, and φ depend on certain fixed hyperparameters.⁴ However, since these hyperparameters are fixed and therefore not random variables, we have chosen not to write them down for clarity. Similarly, all distributions in the joint PPM-EOS inference should be read as conditional on the choices of PPM model (see, e.g., [12,42]) and on the choice of EOS model (see [37,38] for the PP and CS models, respectively). To avoid cluttering the notation, we have omitted including these in the conditional densities as well.

The likelihood is described by

$$\begin{aligned} p(\{D_{k,j}\}_{k,j=1}^{K,J} | \{\theta_k, M_k, \varepsilon_{c,k}\}_{k=1}^K, \varphi) \\ = \prod_{k=1}^K \prod_{j=1}^J p(D_{k,j} | \theta_k, M_k, \varepsilon_{c,k}, \varphi), \end{aligned} \quad (5)$$

⁴Hyperparameters are parameters that determine the shape of the prior distribution for a given model parameter. For example, if the prior for an inclination angle follows a normal distribution, the hyperparameters would be its mean and standard deviation.

where we assume that both neutron stars and data points within each pulse profile are statistically independent. By combining Eqs. (1)–(5), we obtain the posterior probability distribution:

$$\begin{aligned} p(\{\theta_k, M_k, \varepsilon_{c,k}\}_{k=1}^K, \varphi | \{D_{k,j}\}_{k,j=1}^{K,J}) \\ \propto \prod_{k=1}^K \left[\prod_{j=1}^J p(D_{k,j} | \theta_k, M_k, \varepsilon_{c,k}, \varphi) \right] \\ \times p(\theta_k) p(M_k) p(\varepsilon_{c,k}) p(\varphi). \end{aligned} \quad (6)$$

As mentioned before, since the equatorial radius R_k of a given source is completely determined by the combination of $(\varphi, \varepsilon_{c,k}, M_k)$, the posterior distribution on R_k can be computed once the posteriors on φ , $\varepsilon_{c,k}$, and M_k are determined.

B. Intermediate step: PPM including EOS prior information and EOS inferencing

At present, the full hierarchical Bayesian framework described in Sec. II A is computationally intractable. We therefore leave its construction for future work. In this work, we adopt an intermediate approach: first performing PPM with x-PSI, including EOS M – R prior information corresponding to a specific EOS family, and then carrying out the subsequent EOS inference using NEoS. In the PPM stage, we restrict ourselves to a family of EOSs, such that the prior constrains the PPM parameter space to only those combinations permitted within that family. While this builds in physics constraints on the EOS, from chiral effective field theory (χ EFT) calculations at nuclear densities, it also introduces model choices from the high-density extensions. During the EOS inference stage, we then infer which of the parameter combinations within that family are most consistent with the data. This approach has the advantage of being quick to implement while still yielding tighter constraints on neutron star parameters by restricting the parameter space to regions consistent with certain EOSs, while also reducing computational cost.

To construct the EOS-informed posterior distribution for the derived neutron star properties, we begin by formulating the posterior obtained via x-PSI. In previous x-PSI analyses of rotation-powered millisecond pulsars (e.g., [12,14–16,18–20,22]), analysis was done for a single star ($K = 1$) and the EOS prior was not incorporated. Instead, the posterior was given by

$$p(\theta, R, M | \{D_j\}_{j=1}^J) \propto \prod_{j=1}^J p(D_j | \theta, R, M) p(\theta) p(R) p(M), \quad (7)$$

where no explicit EOS model was used to constrain the M – R relation.




FIG. 2. Mass-radius prior distributions for the PP EOS model (left) and CS EOS model (right) when using the $N^3\text{LO } \chi\text{EFT}$ calculations up to $1.5n_0$. The blue contours show the 68%, 95%, and 99.7% credible regions of the EOS prior distributions. Black dotted lines show the corresponding contours from the normalizing flow distributions.

To fold the EOS prior information into the M – R space within x-PSI, we modify Eq. (7) to define a new posterior:

$$p(\theta, R, M | \{D_j\}_{j=1}^J, \varphi) \propto \prod_{j=1}^J p(D_j | \theta, \varphi, R, M) p(\theta) p(R | \varphi) p(M) p(\varphi), \quad (8)$$

where the EOS parameters φ now inform the M – R distribution. In practice, we adopt the following procedure. First, we sample the priors on the EOS parameters and compute their corresponding priors in the M – R space using the N EoST framework (described in more detail below), as shown in Fig. 2. We then use these samples to train a normalizing flow (see Sec. II D) that approximates the EOS M – R prior. Subsequently, if no mass prior is available from radio timing, we draw the mass M from a flat prior. If a mass prior from radio timing is available, we instead sample a mass M from that prior. In either case, we then draw a corresponding radius sample R from the EOS M – R prior conditional on the sampled mass M .

For the EOS inference, we implement the Raaijmakers *et al.* [43] Bayesian framework using N EoST (v2.1.0) [44].⁵ N EoST uses MultiNest to sample from the chosen EOS parameter space and combines constraints from both the M – R relations derived from PPM and the M – Λ relations from gravitational-wave observations. In this work, however, we focus solely on incorporating EOS prior information into the PPM-derived M – R posteriors, and therefore omit the gravitational-wave constraints to highlight the differences between the original and new frameworks. Since the EOS model parameters considered in this work—namely PP and CS—are uniformly sampled, the EOS

inference framework is identical to that used in our previous work [26,43].

The two-step process implemented here naturally raises concerns about double counting the EOS prior, since it is used both to define priors for M – R during the PPM stage and again in the subsequent EOS inference. This is especially relevant because uniform priors on the EOS parameters lead to nonuniform priors on neutron star mass and radius due to the nonlinearity of the TOV equations. However, since these equations are deterministic, the transition from EOS parameters to M – R space corresponds to a change of variables, which is a common operation in Bayesian inference especially when it reduces the computational costs of sampling. Since in our case the prior on the EOS parameters is flat, applying them during both the PPM and subsequently the EOS inference does not distort the overall shape of the EOS posterior, and thus does not influence EOS parameter estimation. It is important to note, however, that Bayesian model comparison via posterior odds should not be performed in this manner, as the inclusion of the prior in both inference steps would indeed distort the results.

Therefore, the posterior distribution of the EOS parameters φ and central densities ϵ_c remains unchanged and is given by

$$p(\varphi, \epsilon_{c,k} | d) \propto p(\varphi) p(\epsilon_{c,k} | \varphi) \prod_{k=1}^K p(M_k, R_k | d_k), \quad (9)$$

where d is the dataset used to constrain the EOS. In this work, d consists of the M – R PPM posteriors for PSR J0740+6620 and PSR J0437-4715 obtained using an EOS-informed M – R prior together with a mass prior derived from radio timing. We do not include other data, such as gravitational-wave observations, in the EOS inference.

It is important to note that when an EOS-informed prior is adopted in the PPM step, the EOS inference is only self-consistent if the same EOS model is used in both PPM and EOS inference. EOS inference with a different EOS model family from that used to define M – R priors for PPM is still possible, provided that the EOS model used to define the PPM priors encompasses the full parameter range of the EOS model applied during the EOS inference. In this scenario, EOS inference is akin to using PPM results derived using EOS-agnostic priors.⁶ This ensures that the

⁵The prerelease form of N EoST was used in [37,43,45–47]. The released version of N EoST has been used in [26,48].

⁶It is important to note that applying different EOS models between the PPM and EOS analysis stages is dependent on whether or not the priors on the EOS model used in the later stage are uniform. If the priors for the EOS model parameters are flat, similar to what is done in this work, then the EOS inference will not be affected by the shape of those priors. However, if the EOS priors are nonuniform, then there is an effective double counting of the prior contribution of the EOS model, thus making the EOS posterior inference statistically inconsistent.

prior informing the PPM-derived $M - R$ distributions—the EOS parameters φ —remains aligned with the model used in the subsequent EOS inference.

C. Pulse profile modeling using x-PSI

We use the open-source x-PSI package, with versions ranging from v3.0.0 to v3.0.5 [31]. As done previously (see, e.g., [16,18–20]), this uses the “Oblate Schwarzschild + Doppler” approximation to model the energy-resolved x-ray pulses from the neutron star [9,10,49–54]. This approximation takes the oblate shape of the neutron star into account and the relativistic effects resulting from the rapid spin of the rotation-powered millisecond pulsars while treating the exterior spacetime as Schwarzschild. For a more detailed description of x-PSI, see [12,31].

x-PSI is used in combination with the nested sampling algorithm MultiNest [55–57] and its Python binding PyMultiNest [58] to compute posterior samples and calculate the evidence for Bayesian model comparison. We opted for this setup over UltraNest [59] due to its lower computational cost, despite UltraNest’s proven convergence and stable results for PSR J0740 + 6620 [22].

For the massive pulsar PSR J0740 + 6620, we use the exact same setup and sampler settings as in the headline result of Salmi *et al.* [20] and Hoogkamer *et al.* [22], where NICER and X-ray Multi-Mirror Mission (XMM-Newton) data are jointly modeled. This setup includes informative priors on mass, inclination, and distance obtained from radio timing [60]. Section IID describes how this mass prior is combined with the EOS-informed $M - R$ prior.

For the nearby and bright pulsar PSR J0437–4715, we likewise adopt the same setup as in the headline result of Choudhury *et al.* [16]. This setup incorporates tight informative priors on mass, inclination, and distance based on radio timing [61] (see Sec. IID for more details), as well as NICER background models to constrain the nonsource background, cross-validated with XMM-Newton data. Instead of 20k live points and a sampling efficiency of 0.3 used by Choudhury *et al.* [16], we use 4k live points and 0.1 sampling efficiency for the EOS-informed prior runs to save on computational costs. These settings are expected to be sufficient, as testing the EOS-agnostic case with 4k live points and 0.1 sampling efficiency yielded results consistent with those obtained using 20k live points and 0.3 sampling efficiency in Choudhury *et al.* [16].

D. Normalizing flow as EOS prior

Uniform sampling in the EOS space induces a nonuniform prior in the $M - R$ space due to the nonlinearity of the TOV equations. Because an analytical transformation of the probability density between these spaces is intractable, we approximate the $M - R$ prior to enable sampling from it during the PPM stage. To achieve this, we draw samples from the uniform prior in EOS space, transform them into corresponding pairs in $M - R$ space by solving the TOV

equations, and then fit a flexible parametric form to the resulting distribution using a normalizing flow.

Normalizing flows (see, e.g., [62–64]) are a class of deep learning models used for density estimation, i.e., learning a function that approximates an unknown probability distribution. Normalizing flows approximate the target distribution by transforming a simple and analytically known distribution, such as a standard normal distribution, into a more complex one by applying a sequence of invertible (i.e., bijective) and differentiable transformations. These transformations are parametrized by neural networks to allow for greater flexibility, and enable both sampling and evaluation of the log-density of the resulting distribution.

In this work, we use normalizing flows to model the EOS $M - R$ prior distribution,

$$p(M, R) = p(R|M)p(M), \quad (10)$$

allowing us to efficiently sample joint $M - R$ combinations, as well as draw radius samples conditioned on a given mass. The latter is particularly useful in cases where a mass prior is available from radio-timing observations, as for PSR J0740 + 6620 and PSR J0437–4715.

In this paper, we use the implementation of normalizing flows provided in Pyro [65],⁷ with a standard normal distribution $\mathcal{N}(\mu = 0, \sigma^2 = 1)$ as the base distribution. For both mass and radius, we use transformations known as coupling neural spline flows [62,66–68], specifically element-wise rational-linear spline bijections, where the radius is conditioned on the mass. Each transformation is parametrized by a small dense feedforward neural network that takes $M - R$ samples from the EOS prior distributions (see Sec. III C) as input. The dataset consists of 100,000 $M - R$ samples of which we reserve 70% for training, 15% for validation, and another 15% for testing. Prior to training, the data are standardized using Pyro’s StandardScaler to have zero mean and unit variance.

To determine the architecture and training hyperparameters,⁸ we perform Bayesian hyperparameter tuning using w&b [69],⁹ aiming to select a configuration that balances validation loss (the negative log-likelihood of the samples given the transformation of the distribution) with training efficiency. The sweep varied the hidden layer dimensions (8–40), the number of transformations (1–10) of $p(R|M)$ and $p(M)$ as shown in Eq. (10), the learning rate ($5 \times 10^{-4} - 5 \times 10^{-2}$), and the number of training epochs (250–12000 in steps of 250).

After performing the Bayesian hyperparameter tuning, we selected the best-performing configuration without further exploring alternative setups. Because there were

⁷<https://pyro.ai>.

⁸In this context, hyperparameters are settings chosen *a priori* that control how the neural network learns, such as the learning rate and the number of hidden layers.

⁹<https://wandb.ai>.




FIG. 3. Training (orange) and validation (blue) loss curves for the PP EOS model (top) and CS EOS model (bottom).

many configurations with similar validation losses, we chose one with a relatively short training time. This configuration—using a normalizing flow with ten transformations for both $p(R|M)$ and $p(M)$ —performed well, achieving a test loss of 2.52 for CS and 2.14 for PP (see Fig. 3 for the training and validation loss curves) in approximately 27 min for both cases. Other configurations with comparable loss values required substantially longer training times. The resulting network contains two hidden layers with 32 neurons each, and ReLU¹⁰ [70] activations following each layer. The normalizing flow is trained on 70,000 M – R samples for 4000 epochs, using the Adam optimizer [71] with a learning rate of 3×10^{-3} . The model is optimized by minimizing the negative log-likelihood of the M – R samples.

To evaluate, we compare the probability distribution learned by the normalizing flow to our test set using the Kullback-Leibler (KL) divergence, which provides a single number quantifying how closely the entire learned distribution approximates the target one. This results in KL divergences of 0.0075 and 0.0023 for the normalizing flow of the CS and PP EOS prior distribution, respectively. A KL divergence close to zero indicates that the distribution learned by the normalizing flow is nearly indistinguishable from the test set. See Fig. 2 for a visual comparison between the prior model and the resultant normalizing flow.

¹⁰ReLU (rectified linear unit) is an activation function defined as $f(x) = \max(0, x)$. Activation functions introduce nonlinearity into the network, allowing the model to capture complex patterns and relationships in the data.

When sampling radii conditioned on a given mass, as in the cases of PSR J0740 + 6620 and PSR J0437-4715, samples with radii outside the interval 5.5 to 14.0 km for the PP EOS model and 6.5 to 14.5 km for the CS EOS model are excluded from the analysis. These bounds are chosen to avoid artifacts arising in regions where the normalizing flow assigns low but nonzero probability and to avoid drawing unphysical M – R values. The bounds are conservative and lie well outside the support of the EOS prior. For a sample size of (1×10^6) , the probability of drawing samples outside these bounds is close to zero, so truncation is not expected to significantly affect the normalization of the distribution.

III. DATA

In this section, we outline the x-ray event data used in this study. In Sec. III A, we describe the most important aspects of the dataset from the high-mass pulsar PSR J0740 + 6620, as detailed in Salmi *et al.* [20] and Hoogkamer *et al.* [22]. Section III B gives a brief description of the observations from the nearest and brightest rotation-powered millisecond pulsar PSR J0437-4715, as detailed in Choudhury *et al.* [16]. In Sec. III C, we summarize the prior distributions of the PP and CS EOS models.

A. PSR J0740 + 6620

For PSR J0740 + 6620, we use the NICER dataset between September 21, 2018, and April 21, 2022, identical to Salmi *et al.* [20] and Hoogkamer *et al.* [22]. We use the pulse-invariant channel subset [30, 150] corresponding to the nominal photon energy range [0.3, 1.5] keV. The data are processed in the same way and include 2.73 Ms (roughly a month) of on-source exposure time after filtering. For background constraints, the same XMM-Newton phase-averaged spectral data and blank-sky observations are used with the three EPIC instruments (pn, MOS1, MOS2), including the energy channels [57, 299] for pn, and [20, 100] for both MOS1 and MOS2 (as reported in [18–20, 22, 72]). Furthermore, the hot emitting regions of the neutron star are modeled with two independent circular uniform temperature spots using the ST-U (single temperature-unshared) model (see [12, 42] for more details on the X-PSI model naming convention and parameters).

B. PSR J0437-4715

Similar to the headline result of PSR J0437-4715 in Choudhury *et al.* [16], we use the 3C50 NICER dataset, which contains observations from July 6, 2017 to October 11, 2021. This dataset owes its name to the usage of the 3C50 background model, which is designed to estimate the instrumental background in NICER observations. The model utilizes NICER event data to predict the background level as a function of the energy (see [73] for the model

definition). To get the best estimate of the instrumental background of NICER, the data are filtered according to specific criteria (see [19] for a detailed description of the analysis and filtering procedure). The filtering steps reduce the usable exposure time from 2.736 to 1.328 Ms, but enable the quantification of the uncertainties in the net background spectrum, which is predominantly influenced by systematic instrumental effects.

To maximize the signal-to-noise ratio and mitigate contamination from the bright Seyfert II active galactic nucleus RX J0437-4711 located $4'.18$ from PSR J0437-4715, NICER observations were conducted pointing $1'.15$ southwest of the pulsar. This deliberate offset reduces the active galactic nuclei (AGN) flux within the field of view to only $\sim 2.25\%$ (i.e., ~ 0.2 counts/s). The AGN flux is then modeled as an additive background component.

For the analysis, only the events in the pulse-invariant channels [30, 300] corresponding to the nominal photon energy range 0.3–30 keV are kept. This is done to match the energy range from XMM-Newton observations, which are used to cross-validate the results from NICER. Lastly, the hot emitting regions of the neutron star are modeled with a two-temperature spot and a ring using the CST + PDT (concentric single temperature + protruding dual temperature model). The two-temperature hot spot consists of two overlapping spherical emitting regions referred to as the ceding and superseding components. The ring is modeled with two spherical components with one emitting and one masking (the latter referred to as omitting in [15,42]).

C. PP and CS EOS prior distributions

With the EOS-informed PPM framework in place (Sec. II B), the next step is to compute the M – R priors for the PP and CS EOS models. This involves constructing the EOS models across the relevant density range and defining the prior ranges that specify their distributions. Following Rutherford *et al.* [26], we briefly summarize the EOS construction adopted in this work and the corresponding priors.

To construct both the PP and CS EOS models, we adopt the BPS crust EOS [35] for densities $\leq 0.5n_0$, with log-linear interpolation to the first χ EFT data points at $0.5792n_0$. For densities in the range $0.5792n_0 \leq n \leq 1.5n_0$, we rely on N³LO χ EFT calculations from Keller *et al.* [36] of neutron star matter in beta equilibrium. Above $1.5n_0$, we extend the EOS to high densities using the PP or CS parametrizations.

The priors on the N³LO χ EFT band are determined by a single polytropic fit, i.e., $P(n) = K(n/n_0)^\Gamma$, to the upper and lower pressure bounds for $n \in [0.5, 1.5]n_0$. This fit yields prior ranges of $K \in [2.207, 3.056]$ MeVfm^{−3} and $\Gamma \in [2.361, 2.814]$ [26].

For the PP model, the priors are defined by varying three polytropic indices, Γ_1 , Γ_2 , and Γ_3 , as well as two transition densities, n_{12} and n_{23} , which connect the first to the second

polytrope and the second to the third, respectively. The priors on these parameters are given by $\Gamma_1 \in [0, 8]$, $\Gamma_2 \in [0, 8]$, $\Gamma_3 \in [0.5, 8]$, and $2n_0 \leq n_{12} \leq n_{23} \leq 8.3n_0$.

For the CS model, the prior distribution is defined by parametrizing the speed of sound $c_s^2 = dP/d\varepsilon$, using five parameters a_1 to a_5 , which are varied over the following ranges: $a_1 \in [0.1, 1.5]$, $a_2 \in [1.5, 12]$, $a_3/a_2 \in [0.05, 2]$, $a_4 \in [1.5, 37]$, and $a_5 \in [0.1, 1]$. The resulting prior distributions for both the PP and CS models span a large range of possible EOSs. For additional details on the PP and CS EOS models, see [26,37,38].

Finally, all EOS priors and posteriors are required to obey causality and the constraint that neutron stars have masses $M \geq 1M_\odot$. This lower mass limit is imposed to be consistent with the lowest confirmed neutron star mass observation [74] and is also motivated by theoretical models of early neutron star evolution [75] as well as the minimum remnant masses predicted by simulations of core-collapse supernovae [76–80]. For completeness, the resulting M – R prior distributions of both of the PP and CS models are shown in Fig. 2.

IV. RESULTS

The results, aimed at showing the impact of incorporating the EOS prior information in the analysis, are presented in two parts. In Sec. IV A, we show the PPM inference results using the EOS M – R prior information for both PSR J0740 + 6620 and PSR J0437-4715. In Sec. IV B, we present the results of the subsequent EOS inference using the posteriors from Sec. IV A.

A. PPM including EOS prior information

To assess the influence of an EOS-informed M – R prior on the inferred properties of PSR J0740 + 6620 and PSR J0437-4715, we evaluate two representative EOS models as priors: CS and PP. These are compared against the baseline EOS-agnostic approach.

1. PSR J0740 + 6620

For PSR J0740 + 6620, we take the headline results of Salmi *et al.* [20], using the EOS-agnostic approach, as our baseline. The analysis of Salmi *et al.* [20], using MultiNest with 4×10^4 live points and a sampling efficiency of 0.01, reported an equatorial radius of $12.49^{+1.28}_{-0.88}$ km and a gravitational mass of $2.073^{+0.069}_{-0.069} M_\odot$. Figure 4 shows the posterior distributions, as well as the ones obtained when incorporating the CS and PP models as prior information on M – R . Table II shows the accompanying KL divergence (in bits) between the posterior and corresponding prior for each model.

The radius credible intervals for all three models overlap. However, the radius inferred using the CS-informed prior is $11.94^{+0.65}_{-0.58}$ km—with the median smaller than the baseline—while the PP-informed prior yielded a slightly




FIG. 4. Radius, compactness, and mass posterior distributions using the PSR J0740 + 6620 joint NICER and XMM-Newton dataset conditional on the ST-U model. Three posterior distributions are shown: the results from Salmi *et al.* [20] without an EOS-informed prior (red), the results obtained with the CS-informed M - R prior (orange), and the ones obtained with the PP-informed M - R prior (blue). All inference results are obtained using MultiNest with 4×10^4 live points and a sampling efficiency of 0.01. The marginal prior probability distribution functions for each parameter are displayed as dashed-dotted lines. The shaded regions in the diagonal panels contain the 68.3% credible interval for each parameter symmetric around the median. The contours in the off-diagonal panels contain the 68.3%, 95.4%, and 99.7% credible regions.

larger median radius of $12.51^{+0.56}_{-0.66}$ km. The results are in line with our expectation, since the CS model favors softer EOSs and the PP model favors stiffer ones. However, once several sources are included, the CS and PP differences are smaller [26]. Notably, the mass and radius credible intervals obtained with the EOS-informed priors are

TABLE II. KL divergence (in bits) between the posterior and corresponding prior for each parameter computed for the three different inferences for PSR J0740 + 6620 shown in Fig. 4.

Parameter	Salmi <i>et al.</i> [20] (EOS agnostic)	CS-informed $M - R$ prior	PP-informed $M - R$ prior
R_{eq} (km)	0.64	0.36	0.11
M/R_{eq}	0.72	0.51	0.14
$M(M_{\odot})$	0.01	0.03	0.01

narrower than those obtained using the EOS-agnostic M - R prior.

The posteriors of the geometry parameters, as shown in Fig. 10 of the Appendix, match even better for both the EOS-agnostic approach and the EOS-informed prior runs, with overlapping credible intervals for all parameters. Furthermore, the difference between the maximum likelihood and Bayesian evidence among all three models is negligible.¹¹

In terms of computational cost, the EOS-agnostic inference required 84k core hours, compared to 21k and 22k core hours for the runs using the CS-informed and PP-informed priors, respectively.

2. PSR J0437-4715

For PSR J0437-4715, we adopted the headline result from Choudhury *et al.* [16] as our baseline, using MultiNest with 20k live points and a sampling efficiency of 0.3. For the EOS-informed prior runs, we used 4×10^3 live points and a sampling efficiency of 0.1 to reduce computational costs. The resulting posterior distributions for all three models are shown in Fig. 5. Table III shows the accompanying KL divergence (in bits) between the posterior and corresponding prior.

As shown in Fig. 5, the EOS-agnostic approach yielded an equatorial radius of $11.36^{+0.95}_{-0.63}$ km and a gravitational mass of $1.418^{+0.037}_{-0.037} M_{\odot}$. The CS-informed prior run inferred a radius of $11.00^{+0.49}_{-0.48}$ km, which is slightly smaller than the baseline result, though the credible intervals otherwise strongly overlap. On the other hand, the PP-informed prior gives a radius of $12.81^{+0.30}_{-0.51}$ km, which is noticeably larger. Similar to PSR J0740 + 6620, these results align with expectations: The CS model favors softer EOSs given a smaller radius, while the PP model favors stiffer ones, giving a larger radius for the same mass. For both EOS modes, the mass and radius credible intervals are again narrower than those obtained from the EOS-agnostic approach.

For both the EOS-informed and EOS-agnostic cases, the inferred geometry consists of a ring surrounding the north pole—the primary hot spot with an omission region in the middle—and a two-temperature spot in the southern hemisphere, the secondary hot spot, consisting of a superseding and ceding region (see Fig. 6 for the inferred maximum-likelihood geometry). Notably, both the CS- and PP-informed prior runs yield geometric solutions that differ in hot spot size, temperature, and location from Choudhury *et al.* [16], as seen in Fig. 7. Most credible intervals for the geometry parameters—especially those of the secondary

¹¹The difference in maximum log-likelihood is approximately 0.1 across all three cases, while the evidence differs by approximately +0.51 for the CS-informed run and +0.24 for the PP-informed run relative to the EOS-agnostic approach.




FIG. 5. Radius, compactness, and mass posterior distributions using the PSR J0437-4715 NICER dataset conditional on the CST + PDT model. Three posterior distributions are shown: the headline result from Choudhury *et al.* [16] using an EOS-agnostic approach (red), the results obtained with the CS-informed M - R prior (orange), and the ones obtained with the PP-informed M - R prior (blue). The EOS-agnostic analysis used MultiNest with 2×10^4 live points and a sampling efficiency of 0.3, while the EOS-informed runs used MultiNest with 4×10^3 live points and a sampling efficiency of 0.1. See Fig. 4 for more details about the figure elements.

hot spot—do not overlap, comparing the EOS-agnostic with the EOS-informed cases. The two EOS-informed runs are dominated by modes that are more extreme than the EOS-agnostic configuration, in terms of their surface emission patterns. In particular, the secondary superseding hot spot appears smaller, hotter, and closer to the south pole, while the ceding component is relatively larger, compared to both the superseding region and the EOS-agnostic solution. The temperature of the ceding

TABLE III. KL divergence (in bits) between the posterior and corresponding prior for each parameter computed for the three different inferences for PSR J0437-4715 shown in Fig. 5.

Parameter	Choudhury <i>et al.</i> [16] (EOS agnostic)	CS-informed $M - R$ prior	PP-informed $M - R$ prior
R_{eq} (km)	1.62	0.59	0.28
M/R_{eq}	1.63	0.66	0.30
$M(M_{\odot})$	0.08	0.07	0.16

component is similar for all cases. On average, the ring is comparable in size to the CS-informed run but becomes larger and thinner (with a larger omission region) in the PP-informed run, relative to the EOS-agnostic solution. For the maximum-likelihood vector, the rings in both the CS- and PP-informed runs are smaller, slightly thicker with a smaller central omission region compared to the EOS-agnostic solution.

Relative to the EOS-agnostic approach, the maximum likelihood and evidence from the CS-informed prior runs are both approximately +38.5 units higher, while those from the PP-informed prior runs are approximately +38.6 and +34.2 units higher, respectively, measured in natural logarithm. These differences in evidence correspond to a Bayes factor of $2 \ln \text{BF} \approx 77.0$ (CS) and 68.4 (PP), constituting very strong evidence in favor of the EOS-informed models, following the interpretation of Kass and Raftery [81]. Furthermore, the CS-informed run is higher than the PP-informed run by +4.3 (evidence), yielding $2 \ln \text{BF} \approx 8.7$ which constitutes strong evidence for the CS-informed model.

Regarding computational cost, the EOS-agnostic prior run required approximately 783k core hours, compared to 33k and 25k core hours for the runs using the CS-informed and PP-informed priors, respectively. Note that these runs are done with different sampler settings, which makes comparing the computational time difficult.

B. EOS inference

We now examine the influence of using the EOS-informed PPM results compared to the EOS-agnostic PPM results for PSR J0740 + 6620 and PSR J0437-4715, as presented in Sec. IV A, on the subsequent EOS inference for the PP and CS models. For the EOS inference only these PPM results are included and no additional astrophysics constraints, such as those from gravitational-wave observations, are used.

Figure 8 shows the pressure-energy density and M - R posteriors for both the PP and CS models obtained with the EOS-agnostic or the EOS-informed PPM results. For the PP model, the EOS-informed PPM results shift the pressure-energy density posterior toward stiffer EOSs, accommodating both the tighter radius constraint for PSR J0740 + 6620 and the larger radii inferred for PSR J0437-4715. Compared to the EOS-agnostic PPM results, the corresponding M - R posteriors yield markedly narrower 68% and 95% credible regions, with an overall preference for larger stellar radii. On the other hand, the CS model with EOS-informed PPM results predicts relatively lower pressures at intermediate densities and correspondingly smaller radii, thereby favoring softer EOSs. Although the EOS-informed PPM results of the PP and CS models prefer stiffer and softer EOSs, respectively, both model posteriors exhibit some overlap. For instance, in the $M - R$ space,




FIG. 6. Projections of the inferred maximum-likelihood geometries are shown for the EOS-agnostic run of Choudhury *et al.* [16] (top) and for the CS-informed $M - R$ prior run (bottom) using the PSR J0437-4715 NICER dataset, both conditional on the CST + PDT model. The projections are shown from the star's equatorial plane (left) and from the observer inclination (right). The color bar shows the normalized hot spot temperature ranging between 0 (coolest, orange) and 1 (hottest, red). In both cases, the geometries consist of a ring surrounding the north pole and a two-temperature spot in the southern hemisphere. Note that in the CS-informed run (bottom), the small (superseding) hot spot in the southern hemisphere lies on the edge of the larger, cooler (ceding) region, whereas in the EOS-agnostic run (top) it is centered within the cooler region.

both EOS-informed PPM results overlap between 12 and 12.6 km. However, when all NICER sources and gravitational-wave constraints are included, the PP and CS posteriors overlap strongly, demonstrating the overall consistency between the two models (Rutherford *et al.* [26]).

Figure 9 shows the posterior pressure distributions at 2, 3, and $4n_0$, illustrating the impact of using EOS-informed PPM results compared to EOS-agnostic PPM results on intermediate-density EOS inferences. For the CS model, the pressures across all three fixed densities shift toward lower central values. On the other hand, for the PP model the pressure shifts toward higher values across all three densities, with narrower and more sharply peaked distributions at $3n_0$ and $4n_0$ when including EOS-informed PPM

results compared to EOS-agnostic PPM results in the EOS inference. Similar to the results in Fig. 8, the EOS-informed PPM pressure posteriors show some overlap across all three fixed densities between the PP and CS models.

V. DISCUSSION AND FUTURE WORK

In this study, we implemented and tested the use of EOS-informed $M - R$ priors in the PPM pipeline and the effect of this on the subsequent EOS inference. We compared this approach to our original EOS-agnostic approach focusing on two pulsars, PSR J0740 + 6620 and PSR J0437-4715, and two EOS models based on χ EFT calculations combined with a CS or PP extension to high densities.




FIG. 7. Radius and spot geometry posterior distributions using the PSR J0437-4715 NICER dataset conditional on the CST + PDT model. Additional details of the parameter descriptions and the prior probability distribution functions can be found in Table A1 of [16] [The 68% credible interval lower limit of the secondary superseding angular radius is -0.0023 , not -0.023 as stated in that Table.]. See Fig. 5 for more details about the figure elements.

A. PPM including EOS prior information

Our main goal was to implement the EOS-informed M - R priors within the PPM framework and, in doing so, obtain tighter constraints on neutron star (geometry) parameters while reducing computational cost. This objective was successfully met: For both pulsars, the credible

intervals derived with EOS-informed priors are narrower than those from the EOS-agnostic approach, and the posteriors behave as expected, with the CS (PP) model favoring somewhat smaller (larger) radii relative to the EOS-agnostic case. As anticipated, the computational costs for PSR J0740 + 6620 were significantly reduced when




FIG. 8. Pressure-energy density (top) and mass-radius (bottom) posteriors for the PP (left) and the CS (right) model when using EOS-agnostic PPM results (dashed/dotted light blue lines/lightly shaded regions) and EOS-informed PPM results (dark blue shaded regions). In the mass-radius space, the EOS-informed PPM results for PSR J0740 + 6620 (orange shaded regions) and PSR J0437-4715 (red shaded regions) are dark shaded, while the EOS-agnostic PPM results are lightly shaded. The dark (light) and inner (outer) areas encompass the 68% (95%) credible regions. The vertical gray lines mark the transition density from the χEFT band to the high-density EOS extensions.

using EOS-informed priors, due to the smaller parameter space the sampler had to explore. For PSR J0437-4715, however, a direct comparison of computational costs could not be made because different sampler settings were used in the EOS-agnostic and EOS-informed runs in order to reduce run-time.

1. PSR J0740 + 6620

For PSR J0740 + 6620, the inferred masses for all three models are almost identical, with the mass posteriors being almost completely dominated by the mass prior from radio timing and the corresponding KL divergences close to zero. This outcome is as expected, given the mass prior from radio timing is highly informative.

Regarding the radius, the credible intervals overlap for all three models. Nevertheless, the median radius obtained with the PP-informed M - R prior resembles that from the EOS-agnostic run of Salmi *et al.* [20] slightly better than the median radius from the CS-informed run. This is expected, since the PP radius prior overlaps more strongly with the EOS-agnostic posterior than the CS prior does, as seen in Fig. 4.

In the case of the PP model, both the radius and the compactness are strongly prior dominated, with KL divergences of 0.11 and 0.14, respectively. Similar behavior is observed, though to a lesser extent, for the CS model, with a KL divergence of 0.36 for the radius and 0.51 for the compactness. Furthermore, for both EOS-informed runs, the posterior distributions extend to the




FIG. 9. Posterior (colored shaded regions) and prior (gray) pressure distributions at $2n_0$ (upper panels), $3n_0$ (middle panels), and $4n_0$ (lower panels). Results are shown for the PP (left) and CS (right) models using the EOS-agnostic (light blue) and EOS-informed (dark blue) PPM results.

upper bound of the radius prior and the lower bound of the compactness prior. Ideally, the physics-informed EOS model (e.g., by the χ EFT constraint) is adapted so that the posterior is shaped primarily by the likelihood, ensuring that the inferred parameters are constrained by the data rather than the prior.

Apart from these differences, no particular model is clearly favored: The maximum likelihood and Bayesian evidence values for all three models are nearly identical across runs.

2. PSR J0437-4715

For PSR J0437-4715, the inferred masses for all three models are again nearly identical, with the mass posteriors largely prior dominated and corresponding KL divergences around 0.1. Similar to PSR J0740 + 6620, this is expected given the highly informative mass prior from radio timing.

Concerning the radius, the credible interval obtained with the CS-informed prior overlaps with the EOS-agnostic result from Choudhury *et al.* [16] at the 68% level, while the PP-informed prior only overlaps at the 95% level, as shown in Fig. 5. In both EOS-informed prior runs, the radius posterior is concentrated around the peak of the corresponding prior. However, for both the radius and the compactness posteriors, the data are clearly informative with KL divergences between 0.28 and 0.66, although the prior still contributes significantly. The EOS-agnostic model yields KL divergences of about 1.6 for both radius

and compactness, indicating a substantial information gain from the data.

It is notable, however, that the CS- and PP-informed $M - R$ prior runs yield significantly different geometric solutions from the one obtained in the EOS-agnostic run of Choudhury *et al.* [16]. In particular, the secondary dual temperature hot region is closer to the south pole in the EOS-informed runs, with the superseding hot spot smaller and hotter than in the EOS-agnostic case. This difference arises because the additional $M - R$ constraints imposed by the EOS models during the PPM stage guide the sampler to explore this region of the parameter space more thoroughly. These solutions are sufficiently well separated in parameter space from those of the EOS-agnostic runs that we can consider them to be a separate mode.

3. Statistical properties of PSR J0437-4715's geometry

Before discussing the physical properties of the newly identified geometric mode, we first discuss its statistical properties. In this regard, the maximum likelihood and evidence values from the CS- and PP-informed prior runs are significantly higher than those obtained in the EOS-agnostic analysis from Choudhury *et al.* [16], with the CS-informed run yielding slightly larger values than the PP-informed run. To determine which mode is formally preferred in a statistical sense, however, requires both modes to be fully explored within the same converged run using the same prior.

To address this, we first revisited the EOS-agnostic run to check whether the new mode had been identified and explored in that analysis. Inspection of the samples from the EOS-agnostic run (from the Zenodo repository [82]), which did not use the multimode exploration setting, shows that MultiNest initially drew samples from regions corresponding to the CS- and PP-informed geometry modes. The sampler even continued to explore this part of parameter space to roughly 150 units below its maximum log-likelihood. However, as the run progressed, the sampler eventually concentrated on the geometric mode reported as the final result in Choudhury *et al.* [16].

In principle, running with the multimode setting on should allow more thorough exploration of multiple different modes by locking in a select number of live points onto identified modes in proportion to the prior mass. In practice, however, this would increase the computational cost further still (the 20k live point run already required 783k core hours, and enabling multimode exploration requires even more live points to properly explore each identified mode). Carrying out such a run is therefore not possible with the available computational resources.

Given these limitations, we instead attempted runs with the more restricted (and hence more computationally

tractable) CS- and PP-informed priors using MultiNest with the multimode exploration setting turned on. However, for the CS-informed prior using 4k live points and a sampling efficiency of 0.3, this resulted in seven distinct modes. One mode resembled the EOS-agnostic solution, while the six remaining modes were similar to each other in terms of most, but not all, parameters of the CS-informed run with the multimode exploration turned off.

By contrast, the PP-informed prior proved more tractable. With 8k live points and a sampling efficiency of 0.3, the multimode run identified two modes: one mode similar to the EOS-agnostic result, and one similar to the EOS-informed prior run, consistent within 95% credible level for all parameters, and within 68% credible level for most parameters. The maximum likelihood and local evidence values for the latter mode, corresponding to the more extreme geometric solution from the EOS-informed prior run, are higher by approximately +35.3 and +29.6 units, respectively, compared to the mode resembling the EOS-agnostic result. This difference in evidence corresponds to a Bayes factor of $2 \ln \text{BF} \approx 59.1$, constituting very strong evidence in favor of the new more extreme geometric mode found by the EOS-informed models, following the interpretation of Kass and Raftery [81].

4. Physical plausibility of PSR J0437-4715's geometry

We now turn to the physical properties of the newly identified geometric mode and assess whether it is physically plausible and consistent with multiwavelength constraints from this source, while noting that more definitive conclusions on its feasibility will require testing against more detailed models and incorporating multiwavelength observational constraints (see, e.g., recent modeling by [83]).

Both the geometric mode obtained in the EOS-agnostic run and the more extreme mode identified in the EOS-informed prior runs feature hot spot geometries composed of combinations of ringlike and circular structures. These configurations can be consistent with the so-called “quadrupolar” magnetic field configuration [84,85]. For instance, the hot spot observed near the northern magnetic pole could plausibly be explained by a nearly aligned dipolar rotator,¹² provided that a fit accommodating an additional smaller spot within the larger ring is viable. In such nearly aligned configurations, the expected current heating pattern includes a compact spot near the magnetic pole—where super-Goldreich-Julian currents are sustained by pair production—and an encircling ring formed by the volume return current [86,87].

¹²Nearly aligned dipolar rotator refers to a neutron star whose magnetic and rotation axes are almost parallel, and whose magnetic field has a simple dipole configuration, implying two antipodal magnetic poles (polar caps).

In contrast, the very small hot (secondary) superseding spot observed in the newly identified mode at the southern pole is hard to reconcile with pulsar electrodynamics. Such narrow, localized structures are problematic, as they support only limited electric potential and are difficult to sustain through pair production processes [88]. In this respect, the somewhat larger superseding spot for the geometric mode identified in Choudhury *et al.* [16] is more physically reasonable.

To further assess whether the inferred hot spot geometry aligns with independent observational constraints, we next consider the viewing geometry of the neutron star as implied by radio observations. The geometrical priors used so far are based on the orbital inclination angle. For a narrow radio emission beam and a fully recycled millisecond pulsar binary, one would expect the orbital inclination angle to be close to the viewing angle between the magnetic axis and the observer's line of sight (e.g., [89]). However, the radio pulse of PSR J0437-4715 is very wide, with multiple connected components that are easily recognizable, spanning at least 200° in pulse longitude. With sufficient sensitivity, low-level radio emission can be detected covering roughly 330° of longitude. This can be explained either by a very wide radio beam, by a nearly aligned rotator (as in the new geometric mode)—which would appear to be inconsistent with the orbital inclination angle information—or by emission originating from a different source.

In addition, we consider the gamma-ray emission as an additional multiwavelength constraint on the inferred geometry. Recently, Kramer and Johnston [90] argued that some radio emission of millisecond pulsars originates not from the polar cap but from beyond the light cylinder, congruent with gamma-ray emission. PSR J0437-4715 indeed shows prominent gamma-ray emission, with significant overlap in pulse phase between the radio and gamma-ray pulse profiles. As argued by Kramer and Johnston [90], this could explain both the wide longitude range of detected radio emission and the complex position angle swing that deviates from a rotating vector model. Solely considering the prominent gamma-ray emission detection suggests a combination of a large viewing angle and magnetic inclination angle, which would be inconsistent with an aligned rotator (see, e.g., [90]). Fitting a rotating vector model is not only extremely difficult, but also inadequate in principle if part of the radio emission does not originate from the polar cap.

A consistent picture is that both the viewing and magnetic inclination angles are large, the radio beam is intrinsically quite wide, and parts of the emission are likely originating from beyond the light cylinder. The geometrical information supplied by the orbital inclination angle is consistent with this image and remains consistent with the adopted priors. All in all, since the newly identified geometric mode in the EOS-informed prior run is close

to a nearly aligned configuration, the less-aligned geometry obtained by the EOS-agnostic run is more plausible according to the gamma-ray and radio observations.

The different, more extreme geometry, also has an effect on the inferred background. The overall median background for the EOS-informed solutions is slightly lower than in the EOS-agnostic case [see the Zenodo repository (footnote 3) for the background plots]. This is consistent with the EOS-informed inferred geometry: The secondary ceding spot is relatively larger, and both the primary and secondary superseding components are (slightly) hotter, contributing more to the unpulsed emission of the pulse profile, which is otherwise compensated by the background in the EOS-agnostic case. Moreover, the secondary hot region (including both ceding and superseding components) lies closer to the south pole. Given the pulsar's viewing angle, this spot remains more visible as the star rotates, further contributing to the unpulsed emission that would otherwise be compensated by the background. Lower and much better constrained backgrounds, as expected from upcoming telescopes eXTP [91] and New-Athena [92], should therefore help in identifying the correct geometric mode.

In conclusion, using EOS-informed priors in the PPM stage enabled the identification of a new mode, whose presence was already suggested by the EOS-agnostic run with multimode exploration setting off, but not fully explored. This mode, which has a more extreme surface emission pattern, is statistically favored over the EOS-agnostic solution by Choudhury *et al.* [16]. Among the EOS-informed prior runs, the CS-informed run is preferred over the PP-informed run according to the Bayes factor, suggesting that the data likely favor smaller radii even for the new geometric mode, consistent with the EOS-agnostic result. Nonetheless, serious questions remain about whether the new, more extreme geometric configuration is physical, in particular with regard to the very small (secondary superseding) hot spot. Future work should address this issue and consider modifying the geometry priors to better reflect what is physically reasonable.

5. Normalizing flow as EOS prior

To model the EOS $M-R$ prior for the PPM step, we used a normalizing flow. We observed that the normalizing flow struggles to capture the $M-R$ distribution near $\approx 2.6M_{\odot}$ and 13 km under the PP model, as shown in Fig. 2. This is likely due to sparse training samples at that boundary, which makes it more difficult for the neural network to learn the distribution edges. However, this is not a concern for our current analysis, since both pulsars have tightly constrained mass priors from radio timing which lie outside this region. We therefore do not sample from this part of the normalizing flow during the PPM step.

B. EOS inference

The use of EOS priors in the PPM stage results in more stringent constraints on the inferred $M-R$ posteriors of PSR J0740 + 6620 and PSR J0437-4715. Including these posteriors in the EOS inference further tightens the PP (CS) EOS posteriors. In the $M-R$ plane, it produces a shift toward larger (smaller) radii and corresponding stiffening (softening) of the pressure-energy density relation for PP (CS), reflecting the expected behavior of these extensions, which relatively favor somewhat stiffer (softer) EOSs during the PPM inference. In particular, PSR J0437-4715 has the strongest impact on this effect.

Studying the impact of EOS-informed PPM results on intermediate-density EOS inferences shows that, for the CS model, the pressures across all three densities consistently shift toward lower values. In contrast, for the PP model the pressures shift to higher values across all three densities. Moreover, at $3n_0$ and $4n_0$, the PP-informed pressure posteriors are more peaked than in the EOS-agnostic case. Overall, pressures increase for the PP model and decrease for the CS model at all densities, consistent with stiffening and softening of the EOS, respectively, when EOS-informed PPM data are used. It is interesting to note that in the EOS analysis with all NICER sources and including gravitational-wave constraints, the prior differences are significantly smaller [26].

C. Future work

In future work, we aim to extend the analysis to sources such as PSR J1231-1411 [23] and PSR J0030 + 0451 [14,15] which exhibit complex, multimodal solutions that are challenging to sample thoroughly within reasonable computational time. A key question is whether incorporating EOS-informed $M-R$ priors during the PPM stage can help mitigate these problems.

Moreover, we note that the current modeling framework is highly dependent on the choice of EOS parametrization. For more general applicability and reduced model bias, it would be advantageous to employ a more agnostic high-density EOS model, for example, based on Gaussian processes [93,94]. While empirical testing would be required to definitively address this, we anticipate that a sufficiently general high-density extension, such as that of a Gaussian process model, would likely converge toward the original EOS-agnostic framework. Nevertheless, the Gaussian process model may retain computational advantages through more efficient prior space constraints compared to uniform sampling over the full causal $M-R$ parameter space. The extent to which this produces tighter posteriors remains uncertain, particularly given that the broad PP and CS priors near $2.0M_{\odot}$ and $1.4M_{\odot}$ still yielded substantially more constrained posteriors compared to the

EOS-agnostic approach for both PSR J0740 + 6620 and PSR J0437-4715, respectively.

This research has made use of data products and software provided by the High Energy Astrophysics Science Archive Research Center, which is a service of the Astrophysics Science Division at NASA/GSFC and the High Energy Astrophysics Division of the Smithsonian Astrophysical Observatory. Software: Cython [95], Daft [96], GetDist [97], GNU Scientific Library [98], HEASoft [99], Matplotlib [100], MPI for Python [101], MultiNest [56], NESTCHECK [102], NumPy [103], PyMultiNest [58], Pyro [65], Python/C language [104], SciPy [105], NEoS [44], W&B [69], and X-PSI [31]. Facilities: NICER, XMM-Newton.

ACKNOWLEDGMENTS

We are grateful to Sasha Philippov for providing input to Sec. V A 2 on the physical implications of the newly identified geometric mode of PSR J0437-4715. M. H. and A. L. W. acknowledge support from the NWO Grant No. ENW-XL OCENW.XL21.XL21.038 Probing the phase diagram of Quantum Chromodynamics (PI: Watts). N. R. acknowledges support by NASA Grant No. 80NSSC22K0092. The work of N. R. was also supported in part by Grant No. NSF PHY-2309135 to the Kavli Institute for Theoretical Physics. A. L. W. and

D. C. acknowledge support from ERC Consolidator Grant No. 865768 AEONS (PI: Watts). The work of M. M., I. S., A. S., and K. H. was supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (Grant Agreement No. 101020842). T. S. acknowledges funding by the Research Council of Finland Grant No. 368807. We acknowledge NWO for providing access to Snellius, hosted by SURF through the Computing Time on National Computer Facilities call for proposals. Part of the work was also carried out on the HELIOS cluster including dedicated nodes funded via the above-mentioned ERC Consolidator grant.

M. H. and N. R. contributed equally to this work and are the joint lead authors.

DATA AVAILABILITY

The data that support the findings of this article are openly available [39].

APPENDIX: SUPPLEMENTARY POSTERIOR PLOT FOR PSR J0740 + 6620

Figure 10 shows the spot geometry and other parameter posterior distributions for PSR J0740 + 6620.




FIG. 10. Spot geometry and other parameter posterior distributions using the PSR J0740 + 6620 NICER and XMM-Newton dataset conditional on the ST-U model. See Fig. 4 for more details about the figure elements.

- [1] A. L. Watts, N. Andersson, D. Chakrabarty, M. Feroci, K. Hebeler, G. Israel, F. K. Lamb, M. C. Miller, S. Morsink, F. Özel, A. Patruno, J. Poutanen, D. Psaltis, A. Schwenk, A. W. Steiner, L. Stella, L. Tolos, and M. van der Klis, *Rev. Mod. Phys.* **88**, 021001 (2016).
- [2] J. M. Lattimer and M. Prakash, *Phys. Rep.* **621**, 127 (2016).
- [3] G. Baym, T. Hatsuda, T. Kojo, P. D. Powell, Y. Song, and T. Takatsuka, *Rep. Prog. Phys.* **81**, 056902 (2018).
- [4] L. Tolos and L. Fabbietti, *Prog. Part. Nucl. Phys.* **112**, 103770 (2020).
- [5] K. C. Gendreau *et al.*, in *Space Telescopes and Instrumentation 2016: Ultraviolet to Gamma Ray*, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9905, edited by J.-W. A. den Herder, T. Takahashi, and M. Bautz (SPIE, Edinburgh, UK, 2016), p. 99051H.
- [6] M. A. Ruderman and P. G. Sutherland, *Astrophys. J.* **196**, 51 (1975).
- [7] J. Arons, *Astrophys. J.* **248**, 1099 (1981).
- [8] A. K. Harding and A. G. Muslimov, *Astrophys. J.* **556**, 987 (2001).
- [9] A. L. Watts, *AIP Conf. Proc.* **2127**, 020008 (2019).
- [10] S. Bogdanov, F. K. Lamb, S. Mahmoodifar, M. C. Miller, S. M. Morsink, T. E. Riley, T. E. Strohmayer, A. K. Tung, A. L. Watts, A. J. Dittmann, D. Chakrabarty, S. Guillot, Z. Arzoumanian, and K. C. Gendreau, *Astrophys. J. Lett.* **887**, L26 (2019).
- [11] S. Bogdanov, A. J. Dittmann, W. C. G. Ho, F. K. Lamb, S. Mahmoodifar, M. C. Miller, S. M. Morsink, T. E. Riley, T. E. Strohmayer, A. L. Watts, D. Choudhury, S. Guillot, A. K. Harding, P. S. Ray, Z. Wadiasingh, M. T. Wolff, C. B. Markwardt, Z. Arzoumanian, and K. C. Gendreau, *Astrophys. J. Lett.* **914**, L15 (2021).
- [12] T. E. Riley, A. L. Watts, S. Bogdanov, P. S. Ray, R. M. Ludlam, S. Guillot, Z. Arzoumanian, C. L. Baker, A. V. Bilous, D. Chakrabarty, K. C. Gendreau, A. K. Harding, W. C. G. Ho, J. M. Lattimer, S. M. Morsink, and T. E. Strohmayer, *Astrophys. J. Lett.* **887**, L21 (2019).
- [13] M. C. Miller *et al.*, *Astrophys. J. Lett.* **887**, L24 (2019).
- [14] T. Salmi, S. Vinciguerra, D. Choudhury, A. L. Watts, W. C. G. Ho, S. Guillot, Y. Kini, B. Dorsman, S. M. Morsink, and S. Bogdanov, *Astrophys. J.* **956**, 138 (2023).
- [15] S. Vinciguerra, T. Salmi, A. L. Watts, D. Choudhury, E. T. Riley, P. S. Ray, S. Bogdanov, Y. Kini, S. Guillot, D. Chakrabarty, W. C. Ho, D. Huppenkothen, S. M. Morsink, Z. Wadiasingh, and M. T. Wolff, *Astrophys. J.* **961**, 62 (2024).
- [16] D. Choudhury, T. Salmi, S. Vinciguerra, T. E. Riley, Y. Kini, A. L. Watts, B. Dorsman, S. Bogdanov, S. Guillot, P. S. Ray, D. J. Reardon, R. A. Remillard, A. V. Bilous, D. Huppenkothen, J. M. Lattimer, N. Rutherford, Z. Arzoumanian, K. C. Gendreau, S. M. Morsink, and W. C. G. Ho, *Astrophys. J. Lett.* **971**, L20 (2024).
- [17] M. C. Miller *et al.*, *Astrophys. J. Lett.* **918**, L28 (2021).
- [18] T. E. Riley *et al.*, *Astrophys. J. Lett.* **918**, L27 (2021).
- [19] T. Salmi, S. Vinciguerra, D. Choudhury, T. E. Riley, A. L. Watts, R. A. Remillard, P. S. Ray, S. Bogdanov, S. Guillot, Z. Arzoumanian, C. Chirenti, A. J. Dittmann, K. C. Gendreau, W. C. G. Ho, M. C. Miller, S. M. Morsink, Z. Wadiasingh, and M. T. Wolff, *Astrophys. J.* **941**, 150 (2022).
- [20] T. Salmi, D. Choudhury, Y. Kini, T. E. Riley, S. Vinciguerra, A. L. Watts, M. T. Wolff, Z. Arzoumanian, S. Bogdanov, D. Chakrabarty, K. Gendreau, S. Guillot, W. C. G. Ho, D. Huppenkothen, R. M. Ludlam, S. M. Morsink, and P. S. Ray, *Astrophys. J.* **974**, 294 (2024).
- [21] A. J. Dittmann, M. C. Miller, F. K. Lamb, I. M. Holt, C. Chirenti, M. T. Wolff, S. Bogdanov, S. Guillot, W. C. G. Ho, S. M. Morsink, Z. Arzoumanian, and K. C. Gendreau, *Astrophys. J.* **974**, 295 (2024).
- [22] M. Hoogkamer, Y. Kini, T. Salmi, A. L. Watts, and J. Buchner, *Phys. Rev. D* **112**, 023008 (2025).
- [23] T. Salmi, J. S. Deneva, P. S. Ray, A. L. Watts, D. Choudhury, Y. Kini, S. Vinciguerra, H. T. Cromartie, M. T. Wolff, Z. Arzoumanian, S. Bogdanov, K. Gendreau, S. Guillot, W. C. G. Ho, S. M. Morsink, I. Cognard, L. Guillemot, G. Theureau, and M. Kerr, *Astrophys. J.* **976**, 58 (2024).
- [24] L. Qi, S. Zheng, J. Zhang, M. Ge, A. Li, S.-N. Zhang, F. Lu, H. Peng, L. Zhang, H. Feng, Z. Zhang, Y. Xu, Z. Li, L. Song, S. Zhang, L. Tao, and W. Ye, *Astrophys. J.* **981**, 99 (2025).
- [25] L. Mauviard *et al.*, [arXiv:2506.14883](https://arxiv.org/abs/2506.14883).
- [26] N. Rutherford, M. Mendes, I. Svensson, A. Schwenk, A. L. Watts, K. Hebeler, J. Keller, C. Prescod-Weinstein, D. Choudhury, G. Raaijmakers, T. Salmi, P. Timmerman, S. Vinciguerra, S. Guillot, and J. M. Lattimer, *Astrophys. J. Lett.* **971**, L19 (2024).
- [27] H. Koehn, T. Wouters, H. Rose, P. T. H. Pang, R. Somasundaram, I. Tews, and T. Dietrich, *Phys. Rev. D* **110**, 103015 (2024).
- [28] C. Huang, L. Tolos, C. Providência, and A. Watts, *Mon. Not. R. Astron. Soc.* **536**, 3262 (2025).
- [29] J. Golomb, I. Legred, K. Chatziioannou, and P. Landry, *Phys. Rev. D* **111**, 023029 (2025).
- [30] J.-J. Li, Y. Tian, and A. Sedrakian, *Phys. Lett. B* **865**, 139501 (2025).
- [31] T. E. Riley, D. Choudhury, T. Salmi, S. Vinciguerra, Y. Kini, B. Dorsman, A. L. Watts, D. Huppenkothen, and S. Guillot, *J. Open Source Software* **8**, 4977 (2023).
- [32] T. E. Riley, G. Raaijmakers, and A. L. Watts, *Mon. Not. R. Astron. Soc.* **478**, 1093 (2018).
- [33] G. Raaijmakers, T. E. Riley, and A. L. Watts, *Mon. Not. R. Astron. Soc.* **478**, 2177 (2018).
- [34] G. Raaijmakers, N. Rutherford, P. Timmerman, T. Salmi, A. L. Watts, C. Prescod-Weinstein, I. Svensson, and M. Mendes, *J. Open Source Software* **10**, 6003 (2025).
- [35] G. Baym, C. Pethick, and P. Sutherland, *Astrophys. J.* **170**, 299 (1971).
- [36] J. Keller, K. Hebeler, and A. Schwenk, *Phys. Rev. Lett.* **130**, 072701 (2023).
- [37] S. K. Greif, G. Raaijmakers, K. Hebeler, A. Schwenk, and A. L. Watts, *Mon. Not. R. Astron. Soc.* **485**, 5363 (2019).
- [38] K. Hebeler, J. M. Lattimer, C. J. Pethick, and A. Schwenk, *Astrophys. J.* **773**, 11 (2013).
- [39] M. Hoogkamer, N. Rutherford, D. Huppenkothen, B. Ricketts, A. Watts, M. Mendes, I. Svensson, A. Schwenk, M. Kramer, K. Hebeler, T. Salmi, and D.

- Choudhury, Reproduction package for “Equation-of-state-informed pulse profile modeling,” (2025), [10.5281/zenodo.17257402](https://zenodo.org/record/17257402).
- [40] R. C. Tolman, *Phys. Rev.* **55**, 364 (1939).
- [41] J. R. Oppenheimer and G. M. Volkoff, *Phys. Rev.* **55**, 374 (1939).
- [42] S. Vinciguerra, T. Salmi, A. L. Watts, D. Choudhury, Y. Kini, and T. E. Riley, *Astrophys. J.* **959**, 55 (2023).
- [43] G. Raaijmakers, S. K. Greif, K. Hebeler, T. Hinderer, S. Nissanke, A. Schwenk, T. E. Riley, A. L. Watts, J. M. Lattimer, and W. C. G. Ho, *Astrophys. J. Lett.* **918**, L29 (2021).
- [44] G. Raaijmakers, N. Rutherford, P. Timmerman, T. Salmi, A. L. Watts, C. Prescod-Weinstein, I. Svensson, and M. Mendes, *J. Open Source Software* **10**, 6003 (2025).
- [45] G. Raaijmakers, T. E. Riley, A. L. Watts, S. K. Greif, S. M. Morsink, K. Hebeler, A. Schwenk, T. Hinderer, S. Nissanke, S. Guillot, Z. Arzumanyan, S. Bogdanov, D. Chakrabarty, K. C. Gendreau, W. C. G. Ho, J. M. Lattimer, R. M. Ludlam, and M. T. Wolff, *Astrophys. J. Lett.* **887**, L22 (2019).
- [46] G. Raaijmakers, S. K. Greif, T. E. Riley, T. Hinderer, K. Hebeler, A. Schwenk, A. L. Watts, S. Nissanke, S. Guillot, J. M. Lattimer, and R. M. Ludlam, *Astrophys. J. Lett.* **893**, L21 (2020).
- [47] N. Rutherford, G. Raaijmakers, C. Prescod-Weinstein, and A. Watts, *Phys. Rev. D* **107**, 103051 (2023).
- [48] N. Rutherford, C. Prescod-Weinstein, and A. Watts, *Phys. Rev. D* **111**, 123034 (2025).
- [49] M. C. Miller and F. K. Lamb, *Astrophys. J. Lett.* **499**, L37 (1998).
- [50] N. R. Nath, T. E. Strohmayer, and J. H. Swank, *Astrophys. J.* **564**, 353 (2002).
- [51] J. Poutanen and M. Gierliński, *Mon. Not. R. Astron. Soc.* **343**, 1301 (2003).
- [52] S. M. Morsink, D. A. Leahy, C. Cadeau, and J. Braga, *Astrophys. J.* **663**, 1244 (2007).
- [53] K. H. Lo, M. C. Miller, S. Bhattacharyya, and F. K. Lamb, *Astrophys. J.* **776**, 19 (2013).
- [54] M. AlGendy and S. M. Morsink, *Astrophys. J.* **791**, 78 (2014).
- [55] F. Feroz and M. P. Hobson, *Mon. Not. R. Astron. Soc.* **384**, 449 (2008).
- [56] F. Feroz, M. P. Hobson, and M. Bridges, *Mon. Not. R. Astron. Soc.* **398**, 1601 (2009).
- [57] F. Feroz, M. P. Hobson, E. Cameron, and A. N. Pettitt, *Open J. Astrophys.* **2**, 10 (2019).
- [58] J. Buchner, PyMultiNest: Python interface for MultiNest, Astrophysics Source Code Library, record ascl:1606.005 (2016).
- [59] J. Buchner, *J. Open Source Software* **6**, 3001 (2021).
- [60] E. Fonseca *et al.*, *Astrophys. J. Lett.* **915**, L12 (2021).
- [61] D. J. Reardon, M. Bailes, R. M. Shannon, C. Flynn, J. Askew, N. D. R. Bhat, Z.-C. Chen, M. Curylo, Y. Feng, G. B. Hobbs, A. Kapur, M. Kerr, X. Liu, R. N. Manchester, R. Mandow, S. Mishra, C. J. Russell, M. Shamohammadi, L. Zhang, and A. Zic, *Astrophys. J. Lett.* **971**, L18 (2024).
- [62] L. Dinh, D. Krueger, and Y. Bengio, [arXiv:1410.8516](https://arxiv.org/abs/1410.8516).
- [63] G. Papamakarios, E. Nalisnick, D. Jimenez Rezende, S. Mohamed, and B. Lakshminarayanan, *J. Mach. Learn. Res.* **22**, 2617 (2021).
- [64] I. Kobyzev, S. J. D. Prince, and M. A. Brubaker, *IEEE Trans. Pattern Anal. Mach. Intell.* **43**, 3964 (2021).
- [65] E. Bingham, J. P. Chen, M. Jankowiak, F. Obermeyer, N. Pradhan, T. Karaletsos, R. Singh, P. Szerlip, P. Horsfall, and N. D. Goodman, [arXiv:1810.09538](https://arxiv.org/abs/1810.09538).
- [66] C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios, [arXiv:1906.04032](https://arxiv.org/abs/1906.04032).
- [67] H. M. Dolatabadi, S. Erfani, and C. Leckie, [arXiv:2001.05168](https://arxiv.org/abs/2001.05168).
- [68] T. Müller, B. McWilliams, F. Rousselle, M. Gross, and J. Novák, [arXiv:1808.03856](https://arxiv.org/abs/1808.03856).
- [69] L. Biewald, Experiment tracking with weights and biases, 2020), software available from <https://www.wandb.com/>.
- [70] V. Nair and G. E. Hinton, in *Proceedings of the 27th International Conference on Machine Learning (ICML-10)* (2010), pp. 807–814.
- [71] D. P. Kingma and J. Ba, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [72] M. T. Wolff, S. Guillot, S. Bogdanov, P. S. Ray, M. Kerr, Z. Arzumanyan, K. C. Gendreau, M. C. Miller, A. J. Dittmann, W. C. G. Ho, L. Guillemot, I. Cognard, G. Theureau, and K. S. Wood, *Astrophys. J. Lett.* **918**, L26 (2021).
- [73] R. A. Remillard, M. Loewenstein, J. F. Steiner, G. Y. Prigozhin, B. LaMarr, T. Enoto, K. C. Gendreau, Z. Arzumanyan, C. Markwardt, A. Basak, A. L. Stevens, P. S. Ray, D. Altamirano, and D. J. K. Buisson, *Astron. J.* **163**, 130 (2022).
- [74] J. G. Martinez, K. Stovall, P. C. C. Freire, J. S. Deneva, F. A. Jenet, M. A. McLaughlin, M. Bagchi, S. D. Bates, and A. Ridolfi, *Astrophys. J.* **812**, 143 (2015).
- [75] K. Strobel, C. Schaab, and M. K. Weigel, *Astron. Astrophys.* **350**, 497 (1999).
- [76] H. T. Janka, B. Müller, F. S. Kitaura, and R. Buras, *Astron. Astrophys.* **485**, 199 (2008).
- [77] T. Fischer, S. C. Whitehouse, A. Mezzacappa, F. K. Thielemann, and M. Liebendörfer, *Astron. Astrophys.* **517**, A80 (2010).
- [78] D. Radice, A. Burrows, D. Vartanyan, M. A. Skinner, and J. C. Dolence, *Astrophys. J.* **850**, 43 (2017).
- [79] Y. Suwa, T. Yoshida, M. Shibata, H. Umeda, and K. Takahashi, *Mon. Not. R. Astron. Soc.* **481**, 3305 (2018).
- [80] B. Müller, A. Heger, and J. Powell, *Phys. Rev. Lett.* **134**, 071403 (2025).
- [81] R. E. Kass and A. E. Raftery, *J. Am. Stat. Assoc.* **90**, 773 (1995).
- [82] Devarshi Choudhury *et al.*, Reproduction package for: “A NICER View of the Nearest and Brightest Millisecond Pulsar: PSR J0437–4715” (2024), [10.5281/zenodo.13766753](https://zenodo.org/record/13766753).
- [83] J. Pétri, S. Guillot, L. Guillemot, D. González-Caniulef, F. Jankowski, J. M. Grießmeier, G. Theureau, and I. Cognard, *Astron. Astrophys.* **701**, A39 (2025).
- [84] S. E. Gralla, A. Lupsasca, and A. Philippov, *Astrophys. J.* **851**, 137 (2017).
- [85] A. Y. Chen, Y. Yuan, and G. Vasilopoulos, *Astrophys. J. Lett.* **893**, L38 (2020).

- [86] A. A. Philippov, A. Spitkovsky, and B. Cerutti, *Astrophys. J. Lett.* **801**, L19 (2015).
- [87] S. E. Gralla, A. Lupasca, and A. Philippov, *Astrophys. J.* **833**, 258 (2016).
- [88] A. Chernoglazov, A. Philippov, and A. Timokhin, *Astrophys. J. Lett.* **974**, L32 (2024).
- [89] M. Kramer *et al.*, *Mon. Not. R. Astron. Soc.* **504**, 2094 (2021).
- [90] M. Kramer and S. Johnston, *arXiv:2510.05778*.
- [91] A. Li *et al.*, *Sci. China Phys. Mech. Astron.* **68**, 119503 (2025).
- [92] M. Cruise, M. Guainazzi, J. Aird, F. J. Carrera, E. Costantini, L. Corrales, T. Dauser, D. Eckert, F. Gastaldello, H. Matsumoto, R. Osten, P.-O. Petrucci, D. Porquet, G. W. Pratt, N. Rea, T. H. Reiprich, A. Simionescu, D. Spiga, and E. Troja, *Nat. Astron.* **9**, 36 (2025).
- [93] P. Landry, R. Essick, and K. Chatziioannou, *Phys. Rev. D* **101**, 123007 (2020).
- [94] S. Ng, I. Legred, L. Suleiman, P. Landry, L. Traylor, and J. Read, *Classical Quantum Gravity* **42**, 205008 (2025).
- [95] S. Behnel, R. Bradshaw, C. Citro, L. Dalcin, D. S. Seljebotn, and K. Smith, *Comput. Sci. Eng.* **13**, 31 (2011).
- [96] D. S. Fulford, D. Foreman-Mackey, and D. W. Hogg, Daft: Beautifully rendered probabilistic graphical models, (2025), <https://docs.daft-pgm.org/en/latest/>.
- [97] A. Lewis, *J. Cosmol. Astropart. Phys.* **08** (2025) 025.
- [98] M. Galassi, J. Davies, J. Theiler, B. Gough, G. Jungman, P. Alken, M. Booth, and F. Rossi, *GNU Scientific Library Reference Manual*, 3rd ed. (Network Theory Ltd., Godalming, 2009).
- [99] Nasa High Energy Astrophysics Science Archive Research Center (Heasarc), HEASoft: Unified Release of FTOOLS and XANADU, Astrophysics Source Code Library (2014), ascl:1408.004.
- [100] J. D. Hunter, *Comput. Sci. Eng.* **9**, 90 (2007).
- [101] L. Dalcín, R. Paz, M. Storti, and J. D'Elía, *J. Parallel Distrib. Comput.* **68**, 655 (2008).
- [102] E. Higson, *J. Open Source Software* **3**, 916 (2018).
- [103] S. van der Walt, S. C. Colbert, and G. Varoquaux, *Comput. Sci. Eng.* **13**, 22 (2011).
- [104] T. E. Oliphant, *Comput. Sci. Eng.* **9**, 10 (2007).
- [105] P. Virtanen *et al.* (SciPy 1.0 Contributors), *Nat. Methods* **17**, 261 (2020).

Correction: An incorrect version of Fig. 8 was used for publication and has now been replaced with the correct version.