Transformers with Physics-informed encodings and Simulation-Based inference for robust Gravitational-Wave detection in Pulsar Timing Array data

Subhajit Dandapat [1]
Alvin J. K. Chua [1 2]

[1]*Department of Physics, National University of Singapore, 21 Lower Kent Ridge Road, Singapore 119077* [2]*Department of Mathematics, National University of Singapore, Singapore 119076*. Correspondence to: Subhajit Dandapat subh_phy@nus.edu.sg.

1. Introduction

Pulsar timing arrays (PTAs) [1–4] provide a unique window into nanohertz gravitational waves (GWs), but extracting astrophysical parameters from noisy, long-baseline timing residuals remains computationally challenging with traditional Bayesian techniques due to high-dimensional parametric spaces, strong degeneracies, and complex, correlated noise models [5–8]. These challenges are particularly acute for deterministic searches such as eccentric supermassive black-hole binaries (ESMBHBs), whose non-sinusoidal, evolving waveforms induce complex, phase-driven timing signatures that are difficult to disentangle using purely data-driven approaches [9, 10].

To address this, we introduce a two-stage physics-informed inference framework, summarized in Fig. 1, built around Transformer architectures [11] with structured positional encodings that explicitly embed analytical GW orbital-phase evolution. In Stage I, a physics-informed Transformer [12] extracts physically meaningful representations from PTA timing residuals and provides accurate point estimates of eccentric binary parameters. In Stage II, the Transformer with physics-informed outputs is summarized into a compact context vector, which conditions discrete or continuous normalizing flows [13, 14], enabling recovery of the full posterior distribution within a simulation-based inference framework [15, 16].

Across a range of signal-to-noise ratio (SNR), the proposed approach achieves sharper posteriors and orders-of-magnitude faster inference than physics-agnostic baselines. The modular framework naturally extends to realistic PTA analyses incorporating red noise and additional signal components, highlighting the promise of physics-aware deep-learning models for next-generation GW inference in PTA- and LISA-like observatories.

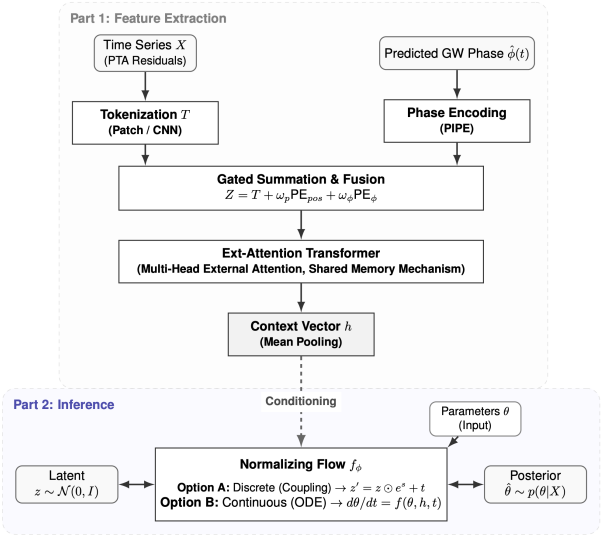

Fig. 1: Overall architecture of the physics-informed SBI pipeline. **Top:** physics-informed Transformer for feature extraction and point estimation. **Bottom:** simulation-based inference via a conditional normalizing flow conditioned on the Transformer context vector h.

2. Dataset Generation

In PTAs, the fundamental observable is the pulsar timing residual. The dataset consists of timing residuals : $X = \{d_j(t_i)\}_{j=1,\ldots,P}^{i=1,\ldots,L}$, measured across P pulsars and L observation times. Each residual is modeled as

$$d_j(t) = s_j(t;\theta) + n_j(t),$$

where $s_j(t;\theta)$ denotes the deterministic GW signal from an ESMBHB system [10], and $n_j(t)$ represents stochastic noise.

We construct simulated GW-induced timing residuals using a standard post-Newtonian signal model [17], in which the residual morphology is governed primarily by the orbital phase $\phi(t)$ evolution. The target parameters are

$$\theta = \{\log_{10} n,\ e_0,\ \log_{10} \mathcal{M}_c,\ \log_{10} S\},$$

corresponding to the orbital frequency, initial eccentricity, chirp mass, and amplitude scale. Noisy realizations are generated by injecting white Gaussian noise and fixing the SNR ranges, enabling controlled evaluation across different noise levels.

3. Physics-Informed Representations

ESMBHB systems exhibit phase-dominated temporal evolution, with $\phi(t)$ governing waveform modulation and periastron advance—structure that index-based positional encodings fail to capture, particularly under irregular sampling or low SNR. To incorporate this domain knowledge, we introduce a *physics-informed phase encoding* (PIPE), where the instantaneous (or predicted) orbital phase is wrapped into $(-\pi, \pi]$, averaged within each token, and mapped to a sinusoidal basis:

$$\mathrm{PE}^{\phi}(b,s,2i) = \sin\left(\frac{\phi^{\mathrm{tok}}_{b,s}}{10000^{2i/d}}\right),$$
$$\mathrm{PE}^{\phi}(b,s,2i+1) = \cos\left(\frac{\phi^{\mathrm{tok}}_{b,s}}{10000^{2i/d}}\right). \quad (1)$$

The resulting embedding tensor $\mathrm{PE}^{\phi} \in \mathbb{R}^{B\times S\times d}$ is combined with standard positional encoding via learnable gates,

$$Z_s = T_s + \omega_{\mathrm{pos}}\mathrm{PE}_{\mathrm{pos}}(s) + \omega_{\phi}\mathrm{PE}^{\phi}(s), \quad (2)$$

allowing the network to adaptively balance physical and positional information.

4. Phase Prediction for Realistic Data

In realistic PTA observations, the true orbital phase is not directly accessible. We therefore train a separate neural network to predict the instantaneous orbital phase from the timing residuals, $x(t) \mapsto \hat{\phi}(t)$, with the phase represented on the unit circle as $(\cos\phi, \sin\phi)$ to respect its periodic nature.

The training objective combines circular alignment, temporal smoothness, and spectral consistency,

$$\mathcal{L}_{\phi} = -\mathbb{E}[\hat{\mathbf{y}}\cdot\mathbf{y}] + \alpha\,\mathbb{E}\left[(\Delta\hat{\phi})^2\right] + \beta\left\||\widetilde{\mathcal{F}[\hat{\mathbf{y}}]}| - |\widetilde{\mathcal{F}[\mathbf{y}]}|\right\|_2^2, \quad (3)$$

where $\mathbf{y} = (\cos\phi, \sin\phi)$ and $\hat{\mathbf{y}}$ denote the true and predicted phase embeddings, respectively.

The predicted phase is subsequently used to construct the physics-informed phase encoding (PIPE), enabling fully data-driven yet phase-aware inference.

5. Two-Stage Inference Framework

5.1 Stage I: Point Estimation with Physics-Informed Transformer

As shown in the **top panel of Fig. 1**, PTA residuals are tokenized and processed by an external-attention Transformer. The sequence output is mean-pooled to obtain a latent representation $h = \frac{1}{S}\sum_{s=1}^{S} Z_s^{(\mathrm{final})}$. A lightweight regression head maps h to point estimates, $\hat{\theta} = f_{\mathrm{MLP}}(h)$, providing fast and accurate parameter predictions. This stage already benefits substantially from physics-informed encoding, particularly for phase-sensitive parameters such as n and e_0.

5.2 Stage II: Simulation-Based Inference with Normalizing Flows

While point estimates are informative, full scientific inference requires the posterior distribution $p(\theta|X)$. In **Stage II**, the Transformer output h—learned using physics-informed phase encoding—serves as a compact *context vector* that summarizes the physically meaningful information extracted from the PTA data (Fig. 1, bottom). This context vector conditions a normalizing flow (NFs), implemented using both discrete and continuous formulations, defining the posterior through an invertible mapping from a simple base distribution

$$z \sim \mathcal{N}(0, I) \;\mapsto\; \theta = f_{\varphi}(z; h). \quad (4)$$

The corresponding posterior density is obtained via the standard change-of-variables formula.

Table 1: Mean log posterior density (LPD) at the true parameters, averaged over 10^4 independent validation realizations. Higher values indicate better-calibrated posterior representations.

Model	$\langle \log q_{\phi}(\theta^{\star} \mid x)\rangle$
CNF (no phase)	−2.196
CNF (predicted phase)	−0.328
DNF (no phase)	−2.509
DNF (predicted phase)	−0.386

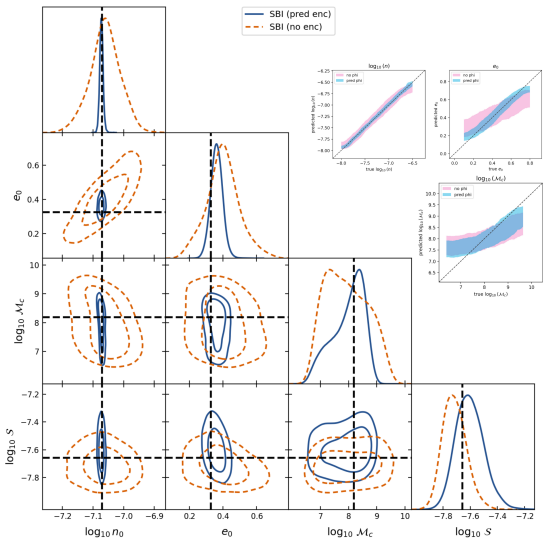

Fig. 2: The top-right panels show point-based inference results (SNR 20–30), comparing predicted and true parameters with and without phase-informed encoding and demonstrating improved calibration when phase information is included. The left panel shows simulation-based inference using a continuous NFs for a single validation realization, where phase-informed encoding (solid blue) yields tighter, better-centered posteriors with reduced parameter degeneracies compared to the encoding-free baseline (dashed orange).

6. Result

Consistent with Fig. 1, physics-informed phase encoding substantially improves both point-based and posterior representations (Fig. 2). Table 1 shows that discrete and continuous normalizing flows conditioned on the predicted phase assign significantly higher log posterior density at the true parameters than encoding-free baselines, indicating improved calibration.

Acknowledgments

This work was supported by the Faculty of Science, National University of Singapore (Grant No. A-8000554-02-00). Computational resources, including GPU allocations, were provided by the National University of Singapore High Performance Computing (NUS HPC) Centre under the CFP03 Project Resource Allocation Award.

References

[1] Gabriella Agazie, Akash Anumarlapudi, Anne M. Archibald, Zaven Arzoumanian, Paul T. Baker, et al. The nanograv 15 yr data set: Evidence for a gravitational-wave background. *The Astrophysical Journal Letters*, 951(1):L8, jun 2023.

[2] J. Antoniadis et al. The second data release from the European Pulsar Timing Array - II. Customised pulsar noise models for spatially correlated gravitational waves. *Astron. Astrophys.*, 678:A49, 2023.

[3] Daniel J. Reardon, Andrew Zic, Ryan M. Shannon, George B. Hobbs, Matthew Bailes, et al. Search for an isotropic gravitational-wave background with the parkes pulsar timing array. *The Astrophysical Journal Letters*, 951(1):L6, jun 2023.

[4] Heng Xu, Siyuan Chen, Yanjun Guo, Jinchen Jiang, Bojun Wang, et al. Searching for the nanohertz stochastic gravitational wave background with the chinese pulsar timing array data release i. *Research in Astronomy and Astrophysics*, 23(7):075024, jun 2023.

[5] Stanislav Babak, Mikel Falxa, Gabriele Franciolini, and Mauro Pieroni. Forecasting the sensitivity of pulsar timing arrays to gravitational wave backgrounds. *Phys. Rev. D*, 110(6):063022, September 2024.

[6] Rutger Van Haasteren, Yuri Levin, Patrick McDonald, and Tingting Lu. On measuring the gravitational-wave background using Pulsar Timing Arrays. *Mon. Not. R. Astron. Soc.*, 395(2):1005–1014, May 2009.

[7] Justin A. Ellis, Michele Vallisneri, Stephen R. Taylor, and Paul T. Baker. Enterprise: Enhanced numerical toolbox enabling a robust pulsar inference suite. Zenodo, September 2020.

[8] Justin Ellis and Rutger van Haasteren. jellis18/ptmcmcsampler: Official release, October 2017.

[9] Gabriella Agazie, Zaven Arzoumanian, Paul T Baker, Bence Bécsy, Laura Blecha, Harsha Blumer, Adam Brazier, Paul R Brook, Sarah Burke-Spolaor, J Andrew Casey-Clyde, et al. The nanograv 12.5 yr data set: a computationally efficient eccentric binary search pipeline and constraints on an eccentric supermassive binary candidate in 3c 66b. *The Astrophysical Journal*, 963(2):144, 2024.

[10] Abhimanyu Susobhanan. Post-newtonian-accurate pulsar timing array signals induced by inspiralling eccentric binaries: accuracy, computational cost, and single-pulsar search. *Classical and Quantum Gravity*, 40(15):155014, jul 2023.

[11] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems (NeurIPS)*, 30, 2017.

[12] Haobo Li, Eunseo Jung, Zixin Chen, Zhaowei Wang, Yueya Wang, Huamin Qu, and Alexis Kai Hon Lau. Pipe: Physics-informed position encoding for alignment of satellite images and time series. *arXiv preprint arXiv:2506.14786*, 2025.

[13] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, 22:1–64, 2021.

[14] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential equations. *Advances in Neural Information Processing Systems (NeurIPS)*, 31, 2018.

[15] George Papamakarios and Iain Murray. Fast ϵ-free inference of simulation models with bayesian conditional density estimation. *Advances in Neural Information Processing Systems (NeurIPS)*, 29, 2017.

[16] Kyle Cranmer, Johann Brehmer, and Gilles Louppe. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 117(48):30055–30062, 2020.

[17] Luc Blanchet. Gravitational radiation from post-newtonian sources and inspiralling compact binaries. *Living reviews in relativity*, 17(1):2, 2014.