

A Physics Informed Bayesian Neural Network for the Neutron Star Equation of State

J.D. Baker,¹ C.A. Bertulani,¹ and R.V. Lobato²

¹*Department of Physics and Astronomy, East Texas A&M University, Commerce, TX 75428, USA*

²*Centro Brasileiro de Pesquisas Físicas, Rua Dr. Xavier Sigaud, 150, Rio de Janeiro, 22290-180, RJ, Brazil*

(Dated: April 29, 2026)

We present a physics-informed Bayesian neural-network framework to infer neutron-star equations of state from theoretical priors and to propagate the associated uncertainties to stellar observables. Trained on a large and representative ensemble of hadronic EoSs, the model learns $P(\epsilon)$ via stochastic variational inference, incorporating soft constraints at saturation density and from perturbative QCD, together with penalties enforcing monotonicity and causality. The accepted core EoSs are matched to an SLy4 crust and evolved through a unified Tolman-Oppenheimer-Volkoff-plus-tidal solver to generate posterior predictions in the mass-radius (M - R) and mass-tidal-deformability (M - Λ) planes. The inferred posterior is consistent with NICER radius measurements and the observed $2.0 M_\odot$ maximum-mass constraint, yielding $R_{1.4} = 12.1_{-0.9}^{+1.4}$ km, $\Lambda_{1.4} = 580_{-240}^{+520}$, and $M_{\text{max}} \simeq 2.11 \pm 0.05 M_\odot$ (90% CI). The resulting canonical tidal deformability can be assessed *a posteriori* against current gravitational-wave constraints. Overall, this framework provides a flexible, non-parametric mapping from microphysical EoS uncertainties to neutron-star observables.

I. INTRODUCTION

The dense-matter equation of state (EoS) remains uncertain above nuclear saturation density, where controlled nuclear-theory constraints and perturbative-Quantum Chromodynamics limits (pQCD) leave a broad intermediate regime. With current multi-messenger measurements, the central task is no longer only to find a single acceptable EoS, but to quantify how microscopic uncertainty propagates to macroscopic observables such as mass, radius, and tidal deformability [1–7].

The modern multi-messenger program was catalyzed by GW170817 and subsequent tidal-deformability analyses, which demonstrated that merger signals can directly constrain the pressure support of matter near and above saturation density [3, 4, 8]. In parallel, Bayesian inference frameworks that combine pulsar masses (massive pulsars [9–11]), X-ray pulse-profile modeling, and gravitational-wave likelihoods have become a standard route to posterior EoS constraints while making prior assumptions explicit [12–14]. More broadly, recent works combine machine-learning surrogates and emulators for neutron-star structure with Bayesian and semiparametric multimessenger EoS inference, including covariant-density-functional, likelihood-systematics, and speed-of-sound-constrained analyses [15–25].

In this work, we introduce a physics-informed Bayesian neural-network framework that performs this uncertainty propagation in a single end-to-end pipeline. Using a large number of hadronic EoSs from the CompOSE database [26], we infer a flexible probabilistic surrogate for $P(\epsilon)$, anchor it with soft constraints at nuclear-saturation and pQCD regimes together with derivative-level monotonicity and causality penalties during Stochastic Variational Inference (SVI) [27], stitch accepted cores to an SLy4 [28] crust, and propagate the resulting ensemble through unified Tolman-Oppenheimer-Volkoff-plus-tidal (TOV-tidal) integration to mass-radius and

tidal-deformability observables [28–30]. In the present analysis, the GW170817 tidal-deformability estimate is used only as an external posterior comparison rather than as a likelihood component of the inference. The central novelty is the combination of functional flexibility with physically admissible posterior families.

Unlike standard Bayesian inference based on low-dimensional parametric families such as piecewise polytropes and spectral expansions [31, 32], our method learns the posterior directly in the function space, reducing the structural bias of predefined parametrizations. This allows localized stiffness variations and nontrivial sound-speed structure to be represented without imposing them *a priori* through segment boundaries, while still reflecting the hadronic training ensemble and the imposed physics-informed priors. Physical admissibility is made explicit at the posterior-sample level through thermodynamic priors, Evidence Lower Bound (ELBO)-level stability and causality penalties, and a permissive post-training sanity screen. Compared with our previous BPR-HE analysis [33], the present framework offers higher functional expressivity and a tighter mapping from microscopic priors to uncertainty bands in mass, radius, and tidal deformability, while remaining aligned with the broader Bayesian multi-messenger inference program [12–14].

In this work, we highlight three points. First, we infer a non-parametric functional posterior over the EoS, rather than restricting the analysis to a low-dimensional parametric family. Second, the nuclear-saturation, pQCD, derivative-level monotonicity, and causality information are embedded directly in the ELBO during training, so physical guidance enters the inference itself rather than arising only from post-hoc filtering. Third, the resulting posterior is propagated in a unified pipeline to both the mass-radius (M - R) and mass-tidal-deformability (M - Λ) planes, enabling joint uncertainty quantification for the two principal observable channels of current multi-messenger neutron-star inference.

Taken together, the PI-BNN framework provides an end-to-end map from microphysical EoS uncertainty to multimessenger observables. By embedding monotonicity, thermodynamic-stability, and causality guidance directly in the variational objective, the model learns posterior EoS families that can be propagated through TOV-tidal calculations to mass, radius, and tidal deformability. In this way, microphysical EoS posteriors are translated directly into macroscopic predictions that are consistent with current NICER constraints and can be compared a posteriori with external GW170817 tidal-deformability estimates, without introducing ad hoc high-density stiffening.

II. THE EQUATION OF STATE

The equation of state (EoS) of nuclear matter encodes the relationship between pressure, density, and composition of strongly interacting matter. In neutron stars, it determines the mass-radius relation, tidal deformabilities measured in gravitational-wave events, and the maximum mass, linking astrophysical observations to the microphysics of QCD in the non-perturbative regime [1, 34].

Despite sustained progress, inferring the EoS remains challenging because the relevant densities span several orders of magnitude, from the crust ($\rho \ll \rho_0$) to the inner core ($\rho \gtrsim 5\rho_0$), where ρ_0 is the saturation density. In this regime, first-principles calculations are limited: chiral effective field theory provides controlled constraints at and below a few times ρ_0 , typically for pure neutron matter and symmetric nuclear matter rather than directly for charge-neutral, beta-equilibrated neutron-star matter, whereas at asymptotically high densities perturbative QCD yields an EoS band, leaving a broad intermediate-density region where modeling assumptions dominate [1, 2, 34].

From the nuclear-physics perspective, major uncertainties arise from the density dependence of the symmetry energy, many-body correlations, and the role of three-nucleon forces, which affect the stiffness of neutron-rich matter. Additional complications include the possible appearance of new degrees of freedom (hyperons, meson condensates, deconfined quarks), the treatment of finite-temperature effects in mergers and supernovae, and the need to enforce thermodynamic stability and causality in any extrapolation.

A further practical issue is that different microscopic approaches provide EoSs in tabulated form with different density grids, compositions, and matching prescriptions between the crust and the core. Community resources such as the CompOSE/CompStar database standardize access to large ensembles of tabulated EoSs and facilitate systematic comparisons and statistical inference [26].

A. Traditional Parametric Representations of the Equation of State

Historically, the mapping of the unknown dense matter equation of state (EoS) has relied heavily on parametric approximations. These methods attempt to bridge the gap between low-density nuclear theory and high-density perturbative QCD by forcing the thermodynamic variables into predefined mathematical structures. While computationally convenient, these rigid functional forms impose artificial constraints on the resulting physics.

a. Piecewise Polytropes: A standard choice is the piecewise polytropic representation introduced for neutron-star inference in Ref. [31]. The stellar interior is partitioned into density intervals, and in each interval the pressure is modeled as (using ϵ as the density variable for notation consistency):

$$P(\epsilon) = K_i \epsilon^{\Gamma_i}, \quad (1)$$

with constants K_i and adiabatic indices Γ_i fixed by continuity and calibration to data. This form is computationally efficient and maps cleanly to macroscopic observables through TOV integration. Its main limitation is structural: even when $P(\epsilon)$ is continuous, higher derivatives can change abruptly at matching densities, so sharp features in the inferred speed of sound may partly reflect the parameterization rather than genuine microphysics.

b. Spectral Expansions: Spectral representations were developed to enforce smooth thermodynamics by expanding an auxiliary EoS quantity (typically the adiabatic index) in a global basis [32]. In practice, one truncates the expansion after a few coefficients, which yields compact inference models and smooth $P(\epsilon)$ and $c_s^2(\epsilon)$ profiles. The tradeoff is basis dependence: low-order truncations can under-resolve localized structures, whereas higher-order truncations may introduce oscillatory behavior unless regularized. As a result, spectral models are often excellent for smooth families of EoSs but can be less robust when strong, localized phase-transition-like features are present.

c. Empirical Parameter Fits (Taylor Expansions): A complementary approach, common in nuclear theory, expands the energy per nucleon around saturation density in terms of empirical coefficients such as K_0 (incompressibility), J (symmetry energy), and L (symmetry-energy slope) [35, 36]. These expansions provide an interpretable connection between microscopic nuclear constraints and macroscopic neutron-star properties. However, their formal control is strongest near ρ_0 ; far above saturation, neglected higher-order terms and possible composition changes can dominate. Consequently, empirical expansions are most reliable in the near-saturation regime and typically require additional high-density priors (e.g., causality and pQCD matching) for inner-core extrapolation.

d. Bayesian Power Regression model with heteroscedastic errors (BPR-HE): In Ref. [33], we

introduced BPR-HE as a Bayesian alternative to segmented piecewise polytrope fits, trained on 65 hadronic EoSs from the LIGO LALSuite infrastructure [37]. The model uses a continuous power-law mean relation,

$$P(\epsilon) = a\epsilon^b, \quad (2)$$

where a sets the pressure scale and b controls effective stiffness. In log-space, the regression is linear, and BPR-HE captures density-dependent spread with a heteroscedastic likelihood, $y_i \sim \mathcal{N}(\alpha + b \ln \epsilon_i, \sigma^2(\epsilon_i))$, using $\sigma(\epsilon) = s_m(\epsilon/\epsilon_{\text{ref}})^m$. This yields positive (log-normal) pressure predictions and propagates uncertainty as an ensemble of admissible EoS curves rather than a single fit.

Compared with piecewise polytropic models, BPR-HE removed segment-boundary artifacts and provided a more consistent Bayesian uncertainty treatment. Its main limitation is that the mean trend is still tied to a fixed analytic form, which motivates the present PI-BNN framework as a non-parametric Bayesian successor that learns a posterior *distribution over functions*.

III. BEYOND PARAMETRIZATION: LEARNING DISTRIBUTIONS OVER FUNCTIONS

Traditional EoS parameterizations are intrinsically *parametric*: they assume a fixed functional form and then infer a limited set of coefficients. This can restrict the allowed thermodynamic behavior and force epistemic uncertainty to be represented only through uncertainties on a few parameters.

In contrast, in this work, a Physics-Informed Bayesian Neural Network (PI-BNN) is used as a non-parametric surrogate for the dense-matter EoS. Instead of selecting one rigid analytic ansatz, the model learns the mapping $P(\epsilon) = f_\theta(\epsilon)$ directly from data, while Bayesian inference over the network weights (via Stochastic Variational Inference, SVI) produces a posterior *distribution over functions* rather than a single best-fit curve.

This unified surrogate perspective offers two practical advantages. First, it enables flexible interpolation and controlled extrapolation across heterogeneous microscopic inputs (hadronic models, near-saturation constraints, and high-density QCD guidance) without hand-crafted segment boundaries. Second, it propagates uncertainty consistently from microphysics to macroscopic observables, yielding posterior predictive EoS bands and robust predictions for masses, radii, and tidal deformabilities.

A. The CompOSE Database

To train the Bayesian Neural Network (BNN) and anchor its predictions in established nuclear physics, we

construct a comprehensive training dataset sourced from the CompOSE (CompStar Online Supernovae Equations of State) database. CompOSE provides a standardized repository of dense matter EoS models derived from various theoretical frameworks, including relativistic mean-field theories and Skyrme interactions.

For this study, we strictly isolate models categorized as purely *hadronic* matter. We explicitly exclude hybrid models or those featuring phase transitions to deconfined quark matter in the training set. This isolation ensures that the BNN learns the fundamental thermodynamic behavior of nucleonic matter up to intermediate densities. The network is therefore unconstrained by hybrid-matter training data, though its learned priors reflect the hadronic training ensemble and should not be interpreted as spontaneously generating first-order phase transitions absent from that distribution. We extract the thermodynamic sequences mapping energy density (ϵ) to pressure (P) for each hadronic model, applying a domain truncation at $\epsilon = 3000 \text{ MeV/fm}^3$ to restrict the training data to physically relevant neutron star core conditions.

a. Data Preprocessing and Balancing: Raw EoS tables frequently exhibit highly irregular sampling frequencies, typically concentrating a massive number of discrete data points in the low-density regime while sparsely sampling the high-density extremes. If fed directly into a neural network, this class imbalance would cause the loss function to disproportionately prioritize fitting the low-density behavior, degrading the model's accuracy and uncertainty quantification in the critical inner core.

To resolve this, we implement a targeted binning and refinement algorithm. The continuous domain of normalized energy density is divided into $N = 30$ uniform histograms. Within each bin, we randomly sample a maximum of $M = 100$ data points. This ensures a balanced, quasi-uniform distribution of training data across the entire density range, helping ensure that the BNN assigns comparable statistical weight across the regimes represented by the hadronic models.

Finally, neural network optimization algorithms, including Stochastic Variational Inference (SVI), are highly sensitive to the scale of input features. Because energy density and pressure in dense matter span several orders of magnitude, we apply a global max normalization. Writing $\epsilon_{\text{max}} = \max(\epsilon)$ and $P_{\text{max}} = \max(P)$, the refined variables are scaled according to

$$\tilde{\epsilon} = \frac{\epsilon}{\epsilon_{\text{max}}}, \quad \tilde{P} = \frac{P}{P_{\text{max}}}. \quad (3)$$

Here $\tilde{\epsilon}, \tilde{P} \in [0, 1]$. For notational simplicity, all subsequent equations are written in the physical variables ϵ and P . When derivative-based quantities are evaluated through the normalized surrogate $\tilde{P} = f_\theta(\tilde{\epsilon})$, the physical sound speed is recovered by the chain rule,

$$c_s^2 = \left(\frac{\partial P}{\partial \epsilon} \right)_S = \frac{P_{\text{max}}}{\epsilon_{\text{max}}} \left(\frac{\partial \tilde{P}}{\partial \tilde{\epsilon}} \right)_S. \quad (4)$$

A microscopic numerical offset (10^{-5}) is added to the normalized arrays to prevent domain errors (such as taking the logarithm of absolute zero) during the calculation of physical constraints and numerical integration later in the pipeline.

B. Artificial Neural Networks

Artificial Neural Networks (ANNs) are highly expressive computational models capable of approximating complex, non-linear functions. In the context of dense matter physics, an ANN can be deployed as a regression tool to map a given energy density ϵ to its corresponding thermodynamic pressure P .

The most fundamental architecture for this continuous mapping is the Feedforward Neural Network, or Multi-Layer Perceptron (MLP). An MLP consists of an input layer, one or more hidden layers, and an output layer. The transformation of data through the l -th hidden layer is given by the affine transformation followed by a non-linear activation function σ :

$$h_l = \sigma(W_l h_{l-1} + b_l) \quad (5)$$

where W_l represents the weight matrix, b_l is the bias vector, and h_{l-1} is the output of the previous layer (with $h_0 = \epsilon$). The final output layer aggregates the highest-level features to produce the predicted pressure P .

The choice of the activation function σ is critical when modeling physical systems. Standard machine learning tasks heavily rely on the Rectified Linear Unit (ReLU). However, ReLU produces a piecewise linear function with a discontinuous first derivative. Because the thermodynamic stability and causality of an EoS depend entirely on its derivative, and because the squared speed of sound is $c_s^2 = (\partial P / \partial \epsilon)_S$, a network using ReLU will generate unphysical, step-like jumps in c_s^2 . To ensure a well-behaved speed of sound across the density domain, we employ the **Softplus** activation function, a smooth approximation of ReLU that yields a continuously differentiable network output.

1. The Deterministic Limitation

In a standard, deterministic ANN, the network parameters $\theta = \{W_l, b_l\}$ are treated as fixed point estimates. Training proceeds by minimizing a loss function (such as Mean Squared Error) over a training dataset using gradient descent algorithms, yielding a single optimal set of weights $\hat{\theta}$.

While a well-trained deterministic ANN can excellently interpolate data within the regime of known hadronic models, it fundamentally lacks the ability to quantify *epistemic* (model) uncertainty. When extrapolating into the ultra-dense inner core of a neutron star where

training data is absent, a deterministic network will output a single, highly confident prediction that is entirely unsupported by underlying physics. To robustly constrain the EoS and extract physically meaningful credible intervals, we must abandon point estimates and transition to a probabilistic framework.

C. Physics-Informed Bayesian Neural Network

A Bayesian Neural Network (BNN) is the probabilistic extension of a standard artificial neural network. In a deterministic network, training returns a single point estimate for each weight and bias. In a BNN, by contrast, the model learns a probability distribution over parameters. This shift from deterministic mapping to probabilistic inference is essential for dense-matter EoS studies, because it enables native quantification of epistemic uncertainty when extrapolating into poorly constrained high-density neutron-star cores.

At the macro level, a BNN has the same structure as a standard multi-layer perceptron: input layer, hidden layers with nonlinear activations, and output layer. The key difference is microscopic: each weight matrix and bias vector is treated as a random variable with an associated distribution. During forward evaluation, parameters are sampled from these distributions, so repeated passes for the same input produce a predictive distribution of outputs. Its mean is used as the central EoS prediction, while its spread quantifies predictive uncertainty.

In Bayesian inference, the prior encodes parameter beliefs before observing data and serves as an important regularizer. A common and effective choice is a zero-mean isotropic Gaussian prior over weights and biases. Centering at zero encodes a preference for small parameters unless the data strongly support otherwise. In practice, this is closely related to L2 regularization (weight decay), which helps stabilize training and reduce overfitting.

The likelihood models how observed data arise from sampled network parameters. For continuous regression tasks such as $\epsilon \mapsto P$, a Gaussian likelihood around the network prediction is standard. Combining prior and likelihood through Bayes' theorem defines the posterior. Because the exact posterior over many network parameters is analytically intractable, we approximate it with SVI, which provides the scalable uncertainty quantification used in our EoS and mass-radius predictions.

The physics-informed Bayesian neural network (PI-BNN) elevates a purely data-driven surrogate by embedding robust microphysical knowledge directly into the Bayesian graphical model. We represent the EoS as a stochastic function $P(\epsilon) = f_\theta(\epsilon)$, where the network parameters θ carry prior uncertainty. Conditioning on both the training dataset $\mathcal{D}$ and a set of theoretical physical constraints $\mathcal{C}$, the augmented posterior can be

written formally as

$$p(\theta|\mathcal{D}, \mathcal{C}) \propto p(\mathcal{D}|\theta)p(\mathcal{C}|\theta)p(\theta) \quad (6)$$

In this structure, $p(\theta)$ is the standard network prior, and $p(\mathcal{D}|\theta)$ is the data likelihood. The crucial addition is $p(\mathcal{C}|\theta)$, which acts as a physics-based weighting factor (equivalently, a pseudo-likelihood for the soft constraints) rather than as an independent dataset likelihood. It scores the physical validity of the network’s predictions in a way that is mathematically consistent with the ELBO-level `pyro.factor` contributions introduced below. This framework keeps the flexibility of neural surrogates while actively suppressing unphysical extrapolations.

1. BNN Architecture

To reconstruct the equation of state (EoS) with robust uncertainty quantification, we model the mapping from energy density ϵ to pressure P using a Bayesian Neural Network (BNN). Let $f_\theta(\epsilon)$ represent the neural network parametrized by weights and biases θ . Unlike standard deterministic networks, we treat θ as random variables to capture epistemic uncertainty.

We define independent Gaussian priors over all weights and biases:

$$p(\theta) = \prod_i \mathcal{N}(\theta_i | 0, \sigma_w^2) \quad (7)$$

where we set the prior scale $\sigma_w = 0.1$. To ensure continuous and differentiable predictions, a strict requirement for calculating the speed of sound, we employ the smooth `Softplus` activation function in the hidden layers, avoiding the discontinuous derivatives associated with `ReLU`.

To approximate the intractable true posterior $p(\theta|\mathcal{D})$, we use Stochastic Variational Inference (SVI). We introduce a parametrized variational distribution $q_\phi(\theta)$ (a mean-field Gaussian) and optimize the variational parameters ϕ by maximizing the Evidence Lower Bound (ELBO):

$$\mathcal{L}_{\text{ELBO}}(\phi) = \mathbb{E}_{q_\phi(\theta)}[\log p(\mathcal{D}|\theta)] - D_{\text{KL}}(q_\phi(\theta) || p(\theta)) \quad (8)$$

where the likelihood $p(\mathcal{D}|\theta)$ assumes Gaussian observation noise around the predictions: $P_i \sim \mathcal{N}(f_\theta(\epsilon_i), \sigma_{\text{obs}}^2)$, with fixed likelihood noise $\sigma_{\text{obs}} = 0.01$.

These architectural choices are physically motivated rather than arbitrary. An MLP is a natural surrogate because the EoS inference problem is a one-dimensional functional mapping, $\epsilon \mapsto P$, so more elaborate architectures designed for images, sequences, or graphs would add complexity without an obvious physical advantage. The smooth `Softplus` nonlinearity is essential because our downstream observables depend on derivatives of the learned function; in particular, computing $c_s^2 = (\partial P / \partial \epsilon)_S$ requires a stable and continuously differentiable surrogate. Finally,

the mean-field Gaussian variational family provides a practical compromise between posterior flexibility and tractability: it scales efficiently under SVI and repeated TOV-tidal evaluations, whereas richer correlated posterior approximations would substantially increase optimization cost and numerical complexity in the present high-dimensional setting.

D. Physics-Informed Soft Constraints

A physical Equation of State must strictly adhere to thermodynamic stability and causality at all densities. Standard Bayesian approaches often rely entirely on highly inefficient post-sampling rejection filters to achieve this. Instead, we embed these physical laws directly into the Bayesian Neural Network as soft constraints during the variational inference process. While a lightweight post-processing filter is retained as a final safety net to guarantee strict mathematical compliance for the numerical TOV integration, the soft constraints guide the variational posterior so effectively that the acceptance rate of the final ensemble is exceptionally high ($\approx 95\%$), bypassing the extreme inefficiency of standard rejection sampling.

In our probabilistic programming framework (`Pyro` [38]), these soft constraints are implemented through `pyro.factor` statements. Rather than enforcing physical admissibility only after sampling, `pyro.factor` adds custom log-density terms directly to the Evidence Lower Bound (ELBO) during Stochastic Variational Inference. This steers the variational posterior toward physically anchored regions of function space before the final numerical safety-net filter is applied, while preserving the flexibility of the neural surrogate.

We enforce three main physics-informed soft-constraint terms during SVI:

1. **Nuclear Saturation Density:** The EoS must pass through the well-constrained empirical saturation point. We apply a Gaussian penalty centered at $\epsilon_{\text{sat}} = 150 \text{ MeV/fm}^3$ with a target pressure $P_{\text{sat}} = 1.5 \text{ MeV/fm}^3$ and variance σ_{sat}^2 .
2. **Perturbative QCD (pQCD) Limit:** To ensure strict theoretical control over the perturbative expansion, we place the high-density anchor at $\epsilon_{\text{pQCD}} = 10,000 \text{ MeV/fm}^3$, corresponding to a baryon density of approximately $n_B \approx 40 n_0$, following the regime identified by Komoltsev and Kurkela as safely perturbative [39]. The target pressure is set to $P_{\text{pQCD}} = 3,333 \text{ MeV/fm}^3$, so that the anchor respects the approach toward the conformal limit $P \approx \epsilon/3$. We implement this anchor through a Gaussian penalty with variance σ_{pQCD}^2 , using a deliberately broad width to retain the residual uncertainty associated with the running coupling at these extreme densities.

3. **Monotonicity Constraint:** Monotonicity is not imposed as a hard rejection rule during training; instead, we add a soft derivative penalty on an auxiliary density grid through `pyro.factor`, which suppresses negative-slope regions while preserving posterior flexibility. Writing the dense auxiliary grid as ϵ_{mono} , the penalty is

$$\mathcal{L}_{\text{mono}} = \lambda_{\text{mono}} \sum_{\epsilon_i \in \epsilon_{\text{mono}}} \text{ReLU}(\delta - c_s^2(\epsilon_i)), \quad (9)$$

where $c_s^2 = (\partial P / \partial \epsilon)_S$ is obtained by automatic differentiation and $\delta = 0.01$ keeps the learned EoS away from marginally unstable slopes.

The saturation and pQCD anchor terms are as important as the monotonicity penalty: the former constrain the normalization of the EoS in empirically controlled low-density and asymptotic high-density regimes, while the latter controls the local slope between those anchors. For the high-density boundary, we intentionally anchor only in a regime where the perturbative expansion is expected to be under firm control, while keeping the Gaussian width broad enough that the pQCD term acts as a soft asymptotic guide rather than as an artificially rigid high-density prior. In addition, we include an analogous soft causality penalty on the same auxiliary grid, proportional to $\text{ReLU}(c_s^2 - 1)$, so that super luminal behavior is disfavored during optimization rather than only after sampling.

Because these contributions enter through `pyro.factor`, the physics-informed objective can be written as an ELBO augmented by anchor and slope penalties,

$$\begin{aligned} \mathcal{J}_{\text{PI}} = \mathcal{L}_{\text{ELBO}} &+ \lambda_{\text{sat}} \log \mathcal{N}(f_{\theta}(\epsilon_{\text{sat}}) \mid P_{\text{sat}}, \sigma_{\text{sat}}^2) \\ &+ \lambda_{\text{pQCD}} \log \mathcal{N}(f_{\theta}(\epsilon_{\text{pQCD}}) \mid P_{\text{pQCD}}, \sigma_{\text{pQCD}}^2) \\ &- \lambda_{\text{mono}} \sum_{\epsilon_j \in \epsilon_{\text{mono}}} \text{ReLU}(\delta - c_s^2(\epsilon_j)) \\ &- \lambda_{\text{causal}} \sum_{\epsilon_j \in \epsilon_{\text{mono}}} \text{ReLU}(c_s^2(\epsilon_j) - 1), \end{aligned} \quad (10)$$

with the optimized training loss given equivalently by $\mathcal{L}_{\text{Total}} = -\mathcal{J}_{\text{PI}}$. In practice, these contributions reshape the variational posterior toward physically admissible regions and implement Bayesian regularization in function space rather than parameter space.

By injecting the penalties during training, the optimizer dynamically penalizes the network whenever it proposes parameter configurations that would generate unstable or acausal EoSs. In practice, this shifts posterior mass away from unphysical regions of function space before sampling begins, drastically improving the acceptance rate of the final ensemble and leaving only a lightweight downstream sanity screen for rare numerical excursions. The net effect is to make physical

admissibility a native part of Bayesian inference rather than a separate post-processing step.

1. Post-Training Consistency Screening

Because thermodynamic stability and causality are already encouraged directly in the ELBO, samples drawn from the optimized variational posterior $q_{\phi}^*(\theta)$ are typically physically admissible before any downstream screening. We therefore use post-training checks only as a permissive numerical consistency step, not as the primary mechanism for enforcing physical constraints.

For every sampled model $f_{\theta^{(k)}}(\epsilon)$, we compute the physical squared speed of sound c_s^2 analytically via auto-differentiation, with the normalization undone through the chain-rule relation above:

$$c_s^2 = \left(\frac{\partial P}{\partial \epsilon} \right)_S = \frac{P_{\text{max}}}{\epsilon_{\text{max}}} \nabla_{\tilde{\epsilon}} \tilde{f}_{\theta^{(k)}}(\tilde{\epsilon}). \quad (11)$$

We then verify that the sampled curve remains within a relaxed tolerance band, $-0.01 \leq c_s^2 \leq 1.05$, over the relevant high-density core domain ($70 \leq \epsilon \leq 2500 \text{ MeV/fm}^3$), and we also apply a mild finite-increment screen requiring $\Delta P > -0.5 \text{ MeV/fm}^3$. These downstream checks remove only rare catastrophic inversions or interpolation artifacts that survive sampling. The dominant suppression of unphysical behavior occurs earlier, inside the ELBO, which substantially improves the acceptance rate of the final stitched ensemble.

2. Crust-Core Stitching

The physics of low-mass neutron stars ($M \lesssim 1.0 M_{\odot}$) is dominated by the outer and inner crust, a regime not directly constrained by the BNN training, which is restricted to supranuclear densities. To model this region, we adopt the tabulated SLy4 equation of state [28] for the crust.

Rather than using stitching as the primary mechanism for imposing monotonicity, we use it to preserve the already softened monotonic behavior of the sampled core EoSs while maintaining an ordered connection to the SLy4 crust. Let $(\epsilon_{\text{crust,max}}, P_{\text{crust,max}})$ denote the terminal point of the SLy4 crust. For each sampled core EoS $f_{\theta}^{(k)}(\epsilon)$ drawn from the BNN posterior, we retain only those points satisfying

$$\epsilon > \epsilon_{\text{crust,max}}, \quad f_{\theta}^{(k)}(\epsilon) > P_{\text{crust,max}}. \quad (12)$$

The final stitched tabulation is then constructed by concatenating the crust segment with this filtered subset

of the core,

$$\begin{aligned} \mathcal{T}_{\text{stitched}} = \mathcal{T}_{\text{SLy4}} \\ \cup \left\{ (\epsilon, f_{\theta}^{(k)}(\epsilon)) : \epsilon > \epsilon_{\text{crust,max}}, \right. \\ \left. f_{\theta}^{(k)}(\epsilon) > P_{\text{crust,max}} \right\}, \end{aligned} \quad (13)$$

and the continuous stitched function $P(\epsilon)$ used in the TOV solver is obtained by monotone interpolation over $\mathcal{T}_{\text{stitched}}$.

This procedure mainly provides an ordered connection between the SLy4 crust and the accepted core segment. Because the final EoS is obtained by monotone interpolation over the combined tabulation, the stitched curve is continuous at the level of the thermodynamic variables. In that sense, it preserves the softened monotonic behavior already encouraged by the ELBO-level stability and causality penalties together with the permissive post-training sanity screen, while avoiding obvious inversions or folds at the matching interface.

We emphasize that this construction enforces continuity at the level of the thermodynamic variables but does not impose higher-order smoothness conditions (e.g., continuity of the speed of sound c_s^2), which is left for future refinement.

E. Relativistic Structure Integration

To map the posterior EoS samples into observable mass-radius and tidal-deformability space, we solve the TOV background equations [29, 30] in geometric units ($G = c = 1$):

$$\frac{dP}{dr} = -\frac{(\epsilon + P)(m + 4\pi r^3 P)}{r^2(1 - 2m/r)}, \quad (14)$$

$$\frac{dm}{dr} = 4\pi r^2 \epsilon. \quad (15)$$

To ensure mathematically consistent background-structure calculations across the wide crust-to-core dynamical range, we represent the stitched EoS in log-log variables, $(\log_{10} P, \log_{10} \epsilon)$. We then use a monotonic Piecewise Cubic Hermite Interpolating Polynomial (PCHIP), which preserves the ordering of tabulated points and avoids artificial overshoot in $(\partial P / \partial \epsilon)_S$ (and therefore in c_s^2). This keeps the interpolated EoS thermodynamically admissible and prevents interpolation-induced super luminal artifacts during stellar-structure integration.

1. Tidal Deformability and Love Number

In addition to the TOV background, we solve the linear even-parity tidal perturbation equations for each

stellar model using the same interpolated EoS and adaptive integrator. The unified solver evolves the background variables together with a first-order radial metric response variable $y(r)$ (introduced by Hinderer [40–42]), so that the stellar mass M , radius R , and surface tidal response are obtained in a single pass for each central condition. In compact form, the tidal-response equation is

$$r \frac{dy}{dr} + y^2 + y F(r) + r^2 Q(r) = 0, \quad y(0) = 2, \quad (16)$$

where $F(r)$ and $Q(r)$ are the standard background functions determined by the TOV solution and the local sound speed $c_s^2 = (\partial P / \partial \epsilon)_S$. In the same geometric units, they are given by

$$F(r) = \frac{1 - 4\pi r^2 [\epsilon - P]}{1 - 2m/r}, \quad (17)$$

and

$$\begin{aligned} Q(r) = \frac{4\pi}{1 - 2m/r} \left[5\epsilon + 9P + \frac{\epsilon + P}{c_s^2} - \frac{3}{2\pi r^2} \right] \\ - 4 \left[\frac{m + 4\pi r^3 P}{r^2(1 - 2m/r)} \right]^2. \end{aligned} \quad (18)$$

Integrating this equation to the stellar surface gives the surface value $y_R \equiv y(R)$ together with the compactness $C = GM/(Rc^2)$, or equivalently $C = M/R$ in geometric units.

From y_R and C , we evaluate the quadrupolar (second-order) tidal Love number k_2 as

$$k_2 = \frac{8C^5}{5} (1 - 2C)^2 \frac{2 + 2C(y_R - 1) - y_R}{\Xi(C, y_R)}, \quad (19)$$

$$\begin{aligned} \Xi(C, y_R) = 2C [6 - 3y_R + 3C(5y_R - 8)] \\ + 4C^3 [13 - 11y_R + C(3y_R - 2) \\ + 2C^2(1 + y_R)] \\ + 3(1 - 2C)^2 [2 - y_R \\ + 2C(y_R - 1)] \ln(1 - 2C). \end{aligned} \quad (20)$$

The corresponding dimensionless tidal deformability is then defined by

$$\Lambda = \frac{2}{3} k_2 C^{-5}. \quad (21)$$

Because the Bayesian posterior produces thousands of EoS realizations, the coupled TOV-tidal ODE system is integrated with an explicit adaptive Runge-Kutta method, with the surface boundary condition $P(R) = 0$ captured through terminal event handling. To manage the computational load of the ensemble, these

integrations are executed in an embarrassingly parallel CPU framework. This unified TOV-tidal integration makes the pipeline a joint inference framework for mass, radius, and tidal deformability, rather than an M - R -only mapping.

IV. NUMERICAL IMPLEMENTATION

The computational pipeline has two numerical components: probabilistic training of the Bayesian neural network (BNN) through Stochastic Variational Inference (SVI) [27], and integration of the coupled stellar-structure and tidal equations with explicit adaptive Runge-Kutta solvers, terminal event handling, and ensemble-parallel CPU execution. The learning stack is implemented in PyTorch [43] and Pyro [38].

Because the BNN parameter space is high dimensional, direct evaluation of the exact posterior $p(\theta \mid \mathcal{D})$ with traditional Markov Chain Monte Carlo (MCMC) methods is computationally impractical. We therefore perform inference with SVI, representing the pressure as a nonlinear function of the energy density, $P = f_\theta(\epsilon)$, and approximating the posterior with a tractable variational family $q_\phi(\theta)$. In practice, we use Pyro’s **AutoNormal** guide, which yields a mean-field Gaussian approximation over the weights and biases.

The variational parameters ϕ are optimized by maximizing the Evidence Lower Bound (ELBO), or equivalently by minimizing the **Trace_ELBO** loss, using the **Adam** optimizer with learning rate 0.01 for exactly 20,000 steps. We then draw $N = 500$ posterior samples for uncertainty propagation.

A key advantage of Pyro is that domain physics can enter the variational objective directly. The soft-constraint terms summarized in Sec. IIID are implemented as **pyro.factor** contributions to the ELBO, which steer the variational posterior toward the low- and high-density anchors while penalizing unstable or acausal behavior.

The squared speed of sound, $c_s^2 = (\partial P / \partial \epsilon)_S$, needed for these derivative-level penalties is evaluated by automatic differentiation rather than finite differences. Using PyTorch’s autograd engine (**torch.autograd**) through the smooth **Softplus** activations allows efficient thermodynamic checks across training batches and posterior samples.

For each accepted EoS, we solve the coupled TOV-tidal system as an initial-value problem with an explicit adaptive Runge-Kutta 5(4) integrator and dynamic step sizing, starting at $r = 10^{-4}$ km. The stitched core-crust EoS is evaluated with the monotonic log-log interpolation defined above, the radial step is capped at $\Delta r_{\text{max}} = 0.2$ km to control near-surface gradients, and the surface boundary condition $P(R) = 0$ is captured through terminal event handling when the pressure reaches the SLy4 vacuum boundary. To manage the posterior ensemble efficiently, these stellar integrations

are distributed across an embarrassingly parallel CPU pool. The resulting compactness and tidal response are then used to compute k_2 and Λ for each stellar model.

A. Network Architecture and Training Hyperparameters

For reproducibility, the Bayesian surrogate is a fully connected Multi-Layer Perceptron (MLP) with one input node, two hidden layers of 64 neurons each, and one output node. The hidden layers use the **Softplus** activation function because its continuous differentiability is strictly required to calculate the smooth first-order derivatives entering the speed-of-sound quantity c_s^2 .

All weights and biases are assigned independent Gaussian priors $\mathcal{N}(0, 0.1)$ in the implementation. Variational training is performed with **Trace_ELBO** Stochastic Variational Inference using the **Adam** optimizer with learning rate 0.01 for exactly 20,000 steps. The likelihood noise is fixed to $\sigma_{\text{obs}} = 0.01$. The resulting optimization history is shown in Fig. 1, where the ELBO rapidly decreases during the initial learning of the low-density EFT data and the soft physics constraints before entering a long, nearly flat asymptotic regime.

The ELBO hyperparameters and anchor locations are fixed as follows: $\lambda_{\text{mono}} = 5000.0$, $\lambda_{\text{causal}} = 12.0$, $\lambda_{\text{sat}} = 5000.0$, $\lambda_{\text{pQCD}} = 5000.0$, $\epsilon_{\text{pQCD}} = 10,000 \text{ MeV/fm}^3$, $P_{\text{pQCD}} = 3333 \text{ MeV/fm}^3$, $\sigma_{\text{sat}} = 1.0 \text{ MeV/fm}^3$, and $\sigma_{\text{pQCD}} = 500.0 \text{ MeV/fm}^3$.

These choices reflect both physics and numerical stability. The tight likelihood noise $\sigma_{\text{obs}} = 0.01$ is chosen so that the network acts nearly as an exact interpolator in the highly reliable low-density chiral-EFT (cEFT) regime. At high density, we deliberately move the pQCD anchor to $\epsilon_{\text{pQCD}} = 10,000 \text{ MeV/fm}^3$ ($n_B \approx 40 n_0$), where the perturbative series is under stricter theoretical control [39], and set $P_{\text{pQCD}} = 3333 \text{ MeV/fm}^3$ to reflect the approach toward the conformal asymptote $P \simeq \epsilon/3$ [44]. The broad choice $\sigma_{\text{pQCD}} = 500.0 \text{ MeV/fm}^3$ then encapsulates the residual uncertainty associated with the running coupling at these extreme densities without forcing an unrealistically narrow asymptotic band. Finally, the strong hierarchy $\lambda_{\text{mono}} = 5000.0 \gg \lambda_{\text{causal}} = 12.0$ is deliberate: monotonicity is enforced through gradients and is therefore especially sensitive to numerical flutter during training, so it requires a stiff barrier-function penalty; causality, by contrast, is a more global high-density property, and a smaller penalty is sufficient to suppress superluminal behavior without artificially over-stiffening the EoS or destabilizing the ELBO loss.

The SVI convergence was confirmed by monitoring the ELBO loss trajectory as shown in Fig. 1. The loss enters a stable asymptotic regime without visible hyperparameter-driven oscillations well before the final




FIG. 1. ELBO convergence history for the PI-BNN training run. The initial steep drop reflects the network rapidly learning the low-density EFT training data together with the imposed soft saturation-density and pQCD constraints. The long, nearly flat tail then provides clear evidence of stable asymptotic convergence of the variational optimization.

training step, indicating that the variational optimization had effectively settled prior to termination.

V. RESULTS AND DISCUSSION

A. Microphysical Posterior and Thermodynamic Stability

The primary objective of the Physics-Informed Bayesian Neural Network (PI-BNN) is to map the epistemic uncertainty of the dense-matter Equation of State (EoS) while enforcing fundamental physical limits. Figure 2 presents the microscopic results of the filtered PI-BNN posterior.

In the left panel, the reconstructed pressure $P(\epsilon)$ is shown alongside the 90% credible interval (CI). At lower densities, the network predictions are tightly constrained by the hadronic training ensemble and the nuclear saturation soft constraint. They are also qualitatively compatible with low-density chiral-EFT guidance, although that comparison remains indirect because the microscopic benchmarks are usually formulated for pure neutron matter or symmetric nuclear matter rather than charge-neutral, beta-equilibrated neutron-star matter [1, 34]. As the energy density increases into the ultra-dense core regime ($\epsilon \gtrsim 1000 \text{ MeV/fm}^3$), the 90% CI naturally widens. This broadening reflects increasing epistemic uncertainty in regions with sparse training data, where the inference is guided primarily by perturbative QCD (pQCD) boundary information.

The right panel of Figure 2 illustrates the impact of our ELBO-level stability and causality penalties. The squared speed of sound, $c_s^2 = (\partial P / \partial \epsilon)_S$, calculated

analytically via PyTorch’s automatic differentiation, is plotted against the energy density. The mean prediction rises to a peak of $c_s^2 \approx 0.8$ at $\epsilon \approx 1800 \text{ MeV/fm}^3$, corresponding very roughly to $n_B \sim 12 n_0$, before softening at higher densities. Across the retained posterior samples, c_s^2 remains below the causality limit $c_s^2 = 1.0$, consistent with thermodynamically admissible and causal EoS behavior. Physically, this indicates that the posterior favors moderately stiff equations of state at intermediate densities, while allowing for softening at higher densities consistent with causality and pQCD guidance.

This density dependence admits a direct physical interpretation. Near and somewhat above saturation density, the posterior remains comparatively soft, consistent with the hadronic training ensemble and the radii inferred for canonical-mass stars. At intermediate densities well above saturation, culminating near $\epsilon \approx 1800 \text{ MeV/fm}^3$ (very roughly $n_B \sim 12 n_0$), the increase in c_s^2 indicates a stiffer response, corresponding to the additional pressure support needed to sustain stars near the observed $2 M_\odot$ threshold. At the highest densities reached in the posterior, the decline of the mean c_s^2 suggests renewed softening rather than indefinite stiffening, indicating that the inferred EoS balances massive-star support with the high-density guidance from pQCD and causality.

Relative to the conformal value $c_s^2 = 1/3$, the posterior does not remain locked to conformal behavior across the full density range. Instead, it rises above $1/3$ at intermediate densities, signaling moderate stiffening above saturation density, and then bends downward at the highest densities reached by the ensemble, indicating partial softening while remaining comfortably below the causal bound.

B. The Mass-Radius and Tidal-Deformability Posterior

By integrating the coupled TOV-tidal equations over the ensemble of causally stable EoS samples, we project the microphysical uncertainties into the observable planes shown in Figure 3. The left panel displays the mass-radius (M - R) relation, while the right panel shows the corresponding mass-tidal-deformability (M - Λ) relation. To avoid survivorship bias, the 90% Credible Interval (CI) bands are strictly truncated at the median maximum mass of the posterior, preventing statistical artifacts associated with the collapse of softer EoSs at lower masses.

Mass-Radius Plane: The left panel shows that the PI-BNN posterior remains compatible with the $2.0 M_\odot$ limit without any ad hoc high-density stiffening. The 90% Credible Interval overlaps with the NICER credible regions for PSR J0030+0451, PSR J0740+6620, and PSR J0437-4715. The NICER measurements are shown for comparison and are not directly incorporated




FIG. 2. Microphysical predictions of the filtered PI-BNN posterior. Left: pressure as a function of energy density, where the solid blue curve marks the posterior mean and the shaded blue band denotes the 90% credible interval. Right: the corresponding squared speed of sound c_s^2 , where the solid blue curve gives the mean prediction and the red dashed line marks the causality limit $c_s^2 = 1$.

as likelihood constraints in the present framework. This overlap indicates that the same posterior family is compatible with both canonical-mass and near-maximum-mass stars. The high-mass turning region remains narrow, with an estimated $M_{\text{max}} \simeq 2.11 \pm 0.05 M_{\odot}$ (90% CI), while the truncation at the median posterior maximum mass prevents the apparent narrowing of the band from being biased by the earlier termination of softer EoSs [45].

By taking a horizontal slice of our interpolated posterior grid at $M = 1.4 M_{\odot}$, we extract a canonical radius of $R_{1.4} = 12.1^{+1.4}_{-0.9}$ km (90% CI).

Tidal-Deformability Plane: The right panel shows that tidal deformability is a central posterior observable rather than a secondary by-product of the stellar-structure calculation. The posterior displays the expected decrease of Λ with increasing mass, reflecting the greater compactness of heavier stars, and the shape of the 90% Credible Interval indicates that the model avoids unphysical stiffening at intermediate densities. The GW170817 constraint is not used as a likelihood in the present inference, but serves as an external benchmark for the posterior. At canonical mass, our model predicts $\Lambda_{1.4} = 580^{+520}_{-240}$ (90% CI), whereas the GW170817 reference value is $\Lambda_{1.4} = 190^{+390}_{-120}$ (90% CI). The two posteriors are therefore only marginally consistent where their 90% credible intervals overlap. The higher central value obtained here is a natural manifestation of the well-known tension in the literature: because the inferred EoS successfully supports observed $2.0 M_{\odot}$ pulsars, which require a comparatively stiff equation of state, $\Lambda_{1.4}$ is correspondingly pushed upward. Because the GW170817

constraint is not incorporated as a likelihood input here, this offset should be interpreted as an external comparison rather than as a fit discrepancy. Lower-mass configurations populate the larger- Λ region, while more massive stars shift toward smaller deformabilities, consistent with the usual correlation between tidal deformability and radius at fixed mass. In this sense, the physics-informed constraints allow the model to balance this tension naturally without violating either extreme.

VI. CONCLUSIONS

In this work, we introduced a physics-informed Bayesian Neural Network (PI-BNN) to infer the dense matter equation of state and systematically propagate microscopic uncertainties into macroscopic mass-radius and tidal-deformability observables. By treating the EoS as a continuous probability distribution over functions rather than a rigid parametric curve, the method reduces derivative artifacts and structural bias associated with traditional piecewise polytropic and spectral representations while retaining fully Bayesian function-space inference.

By anchoring the neural network with soft constraints at nuclear saturation and perturbative QCD densities, and adding derivative-level monotonicity and causality penalties during SVI, we steer the posterior toward thermodynamically reasonable behavior without relying on post-sampling rejection as the primary admissibility mechanism. A permissive post-training sanity screen is retained only to remove catastrophic numerical




FIG. 3. Joint observable-plane predictions derived from the PI-BNN posterior. Left: the mass-radius (M - R) plane, where the solid blue curve is the posterior mean and the blue shading denotes the 90% credible interval, truncated at the median posterior maximum mass to avoid survivorship bias from the collapse of softer EoSs. The red dashed line marks the $2.0 M_\odot$ limit, and the purple, orange, and crimson patches indicate NICER measurements for PSR J0030+0451, PSR J0740+6620, and PSR J0437-4715, respectively. Right: the mass-tidal-deformability (M - Λ) plane, showing the same posterior mean and truncated credible region together with the external GW170817 tidal-deformability constraint, which is used here for comparison rather than as a likelihood input. The band implies $\Lambda_{1.4} = 580^{+520}_{-240}$ (90% CI), to be compared with the GW170817 value $\Lambda_{1.4} = 190^{+390}_{-120}$ (90% CI); the two are only marginally consistent where their 90% credible intervals overlap, reflecting the familiar tension between lower tidal deformability and the stiffness required to support $2.0 M_\odot$ pulsars.

excursions. The speed-of-sound analysis shows non-linear stiffening without violating the $c_s^2 \leq 1.0$ limit, and the BNN-to-SLY4 stitching with monotonic log-log interpolation yields stable TOV solutions across both low-mass and high-mass stars.

The resulting posterior is consistent with current multi-messenger astrophysics in both observable planes. At low to intermediate densities it is also qualitatively compatible with microscopic chiral-EFT expectations, although that comparison is indirect because such calculations are commonly formulated for symmetric nuclear matter or pure neutron matter, whereas the neutron-star EoS relevant here describes charge-neutral, beta-equilibrated matter [1, 34]. In the mass-radius (M - R) plane, the PI-BNN posterior remains compatible with the $2.0 M_\odot$ maximum-mass constraint and overlaps with NICER credible regions from X-ray pulse profiling, including the refined J0437-4715 measurement. In the mass-tidal-deformability (M - Λ) plane, the posterior yields $\Lambda_{1.4} = 580^{+520}_{-240}$ (90% CI), whereas GW170817 gives $\Lambda_{1.4} = 190^{+390}_{-120}$ (90% CI), so the two are only marginally consistent where their 90% credible intervals overlap. The higher central value in our posterior reflects the familiar tension between the relatively low GW170817 deformability and the stiffness required to support observed $2.0 M_\odot$ pulsars; the present physics-

informed constraints allow the model to balance this tension without violating either extreme. Because that constraint enters only as an external comparison, not as a likelihood input, tidal deformability remains a direct posterior output rather than a fitted diagnostic. More importantly, the framework provides a scalable, non-parametric route for incorporating future multi-messenger constraints directly into EoS inference, so that improved X-ray, radio, and gravitational-wave measurements can be translated transparently into tighter constraints on dense-matter microphysics and neutron-star structure.

In present analysis the training set is restricted to hadronic EoSs; explicit phase transitions and hybrid-matter EoSs are not included. This scope was chosen to establish a controlled baseline with a uniform microphysical dataset and to validate end-to-end uncertainty propagation before introducing additional transition parameters. Extending the framework to include phase transitions and hybrid-matter families is a natural next step and will require dedicated priors and expanded training data.

Acknowledgments

This work was supported by U.S. Department of Energy Office of Nuclear Physics under Contract No.

DE-SC0026074 with East Texas A&M University. RVL was supported by INCT-FNA (Instituto Nacional de

Ciência e Tecnologia, Física Nuclear e Aplicações), research Project No. 464898/2014-5, and acknowledges support from CAPES/CNPq.

-
- [1] M. Oertel, M. Hempel, T. Klähn, and S. Typel, *Reviews of Modern Physics* **89**, 015007 (2017), 1610.03361.
 - [2] O. Komoltsev, R. Somasundaram, T. Gorda, A. Kurkela, J. Margueron, and I. Tews, *Physical Review D* **109**, 094030 (2024), 2312.14127.
 - [3] B. Abbott *et al.* (LIGO Scientific, Virgo), *Phys. Rev. Lett.* **119**, 161101 (2017), 1710.05832.
 - [4] B. Abbott *et al.* (LIGO Scientific, Virgo), *Physical Review Letters* **121**, 161101 (2018), 1805.11581.
 - [5] T. E. Riley *et al.*, *The Astrophysical Journal* **887**, L21 (2019), 1912.05702.
 - [6] T. E. Riley *et al.*, *The Astrophysical Journal Letters* **918**, L27 (2021), 2105.06980.
 - [7] D. Choudhury *et al.*, *The Astrophysical Journal Letters* **971**, L20 (2024), 2407.06789.
 - [8] E. Annala *et al.*, *Phys. Rev. Lett.* **120**, 172703 (2018), 1711.02644.
 - [9] J. Antoniadis *et al.*, *Science* **340**, 6131 (2013), 1304.6875.
 - [10] H. T. Cromartie *et al.* (NANOGrav), *Nature Astron.* **4**, 72 (2019), 1904.06759.
 - [11] E. Fonseca *et al.*, *Astrophys. J. Lett.* **915**, L12 (2021), 2104.00880.
 - [12] A. W. Steiner, J. M. Lattimer, and E. F. Brown, *The Astrophysical Journal* **722**, 33 (2010), 1005.0811.
 - [13] P. Landry, R. Essick, and K. Chatziioannou, *Physical Review D* **101**, 123007 (2020), 2003.04880.
 - [14] G. Raaijmakers *et al.*, *The Astrophysical Journal Letters* **918**, L29 (2021), 2105.06981.
 - [15] K. Zhou *et al.*, *Progress in Particle and Nuclear Physics* **135**, 104084 (2024), 2303.15136.
 - [16] G. Papigkiotis *et al.*, *Physical Review D* **111**, 083056 (2025), 2501.18544.
 - [17] I. Liodis, E. Smirniotis, and N. Stergioulas, *Physical Review D* **109**, 104008 (2024), 2309.03991.
 - [18] Y. Fujimoto *et al.*, *Physical Review D* **110**, 034035 (2024), 2401.12688.
 - [19] N. Stergioulas, in *Compact Objects in the Universe*, edited by E. Papantonopoulos and N. Mavromatos (Cham, 2024) pp. 329–356.
 - [20] S. Ng *et al.*, *Classical and Quantum Gravity* **42**, 205008 (2025), 2507.03232.
 - [21] J.-J. Li, Y. Tian, and A. Sedrakian, *Physical Review C* **111**, 055804 (2025), 2502.20000.
 - [22] X. Dong *et al.*, *Physical Review D* **112**, 063043 (2025), 2509.03069.
 - [23] J.-J. Li and A. Sedrakian, *Physical Review C* **112**, 015802 (2025), 2505.00911.
 - [24] J.-J. Li, Y. Tian, and A. Sedrakian, *Physics Letters B* **865**, 139501 (2025), 2412.16513.
 - [25] G. Papigkiotis, G. Vardakas, and N. Stergioulas, *Physical Review D* **113**, 023018 (2026), 2508.05850.
 - [26] S. Typel *et al.* (CompOSE Core Team), **58**, 221 (2022), 2203.03209.
 - [27] M. D. Hoffman *et al.*, *Journal of Machine Learning Research* **14**, 1303 (2013).
 - [28] F. Douchin and P. Haensel, *Astronomy & Astrophysics* **380**, 151 (2001), astro-ph/0111092.
 - [29] R. C. Tolman, *Physical Review* **55**, 364 (1939).
 - [30] J. R. Oppenheimer and G. M. Volkoff, *Physical Review* **55**, 374 (1939).
 - [31] J. S. Read, B. D. Lackey, B. J. Owen, and J. L. Friedman, *Physical Review D* **79**, 124032 (2009), 0812.2163.
 - [32] L. Lindblom, *Physical Review D* **82**, 103011 (2010), 1009.0738.
 - [33] E. V. Chimsanski, R. V. Lobato, A. R. Goncalves, and C. A. Bertulani, *Particles* **6**, 198 (2023), 2205.01174.
 - [34] K. Chatziioannou, H. T. Cromartie, S. Gandolfi, I. Tews, D. Radice, A. W. Steiner, and A. L. Watts, *Reviews of Modern Physics* **97**, 045007 (2025), 2407.11153.
 - [35] J. M. Lattimer and Y. Lim, *The Astrophysical Journal* **771**, 51 (2013), 1203.4286.
 - [36] J. Margueron, R. Hoffmann Casali, and F. Gulminelli, *Physical Review C* **97**, 025805 (2018), 1708.06894.
 - [37] LIGO Scientific Collaboration, “Ligo algorithm library - lalsuite,” (2018).
 - [38] E. Bingham *et al.*, *Journal of Machine Learning Research* **20**, 1 (2019).
 - [39] O. Komoltsev and A. Kurkela, *Physical Review Letters* **128**, 202701 (2022), 2111.05350.
 - [40] T. Hinderer, *The Astrophysical Journal* **677**, 1216 (2008), 0711.2420.
 - [41] T. Hinderer, B. D. Lackey, R. N. Lang, and J. S. Read, *Physical Review D* **81**, 123016 (2010), 0911.3535.
 - [42] S. Postnikov, M. Prakash, and J. M. Lattimer, *Physical Review D* **82**, 024016 (2010), 1004.5098.
 - [43] A. Paszke *et al.*, (2019), 10.48550/arXiv.1912.01703, 1912.01703.
 - [44] E. Annala *et al.*, *Nature Physics* , 1 (2020), 1903.09121.
 - [45] The low-mass regime ($M \lesssim 0.5M_{\odot}$) exhibits a sharp expansion in radius, characteristic of stellar structures dominated by the nuclear envelope. This behavior is consistent with the adopted crust-core stitching procedure and indicates that low-mass configurations are not artificially truncated. For the sake of simplicity we truncated the plot at $0.5M_{\odot}$ since this region is not of the interest of the Bayesian framework.