

Physics-informed Bayesian neural network for the neutron star equation of state

J. D. Baker¹, C. A. Bertulani¹, and R. V. Lobato²

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

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



(Received 27 April 2026; accepted 9 September 2026; published 29 September 2026)

We present a physics-informed Bayesian neural-network framework for inferring neutron-star equations of state (EOS) from theoretical priors and propagating the resulting uncertainty to stellar observables. Trained on a representative set of hadronic EOS, the model learns the equation of state through stochastic variational inference by representing the squared speed of sound with a bounded network output and obtaining the pressure by integration, so that causality, thermodynamic stability, and monotonicity are guaranteed by construction, with low-density nuclear and perturbative-QCD normalization anchors. Core EOS are matched to a SLy4 crust and propagated through a unified Tolman-Oppenheimer-Volkoff-plus-tidal solver to obtain posterior predictions in the mass-radius (M - R) and mass-tidal-deformability (M - Λ) planes. The physics-informed prior is then updated with current multimessenger data: NICER radius measurements, the GW170817 tidal-deformability constraint, and the $2M_{\odot}$ maximum-mass bound, included directly in the variational objective. The observational update moves the canonical radius from $R_{1.4} = 13.41$ km in the prior to $R_{1.4} = 12.74^{+0.97}_{-0.73}$ km (nominal 90% variational credible interval), with $\Lambda_{1.4} = 428^{+249}_{-130}$ and $M_{\max} \gtrsim 2.0M_{\odot}$. This framework provides a nonparametric route from microphysical EOS uncertainty to neutron-star observables.

DOI: [10.1103/x9r8-wbf7](https://doi.org/10.1103/x9r8-wbf7)

I. INTRODUCTION

The dense-matter equation of state (EOS) remains uncertain above nuclear saturation density, where controlled nuclear-theory constraints and perturbative-QCD limits leave a broad intermediate regime. With current multimessenger 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 multimessenger 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 work combines 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].

Neural networks have by now been applied to essentially every stage of this inference chain. Deterministic networks

have been trained to map mass-radius (or spectral) data to EOS parameters or pressure values [26–34], to reconstruct the EOS via automatic differentiation through the stellar-structure equations [35,36], and to classify EOS families from observables [37]. Nonparametric neural representations of the EOS itself have been sampled within fully Bayesian multimessenger analyses [38,39], Bayesian neural networks have been used to decode neutron-star composition [40], and uncertainty-aware and simulation-based-inference pipelines have been developed to extract EOS information directly from telescope spectra [41–44]. Relative to this body of work, the present framework combines two inference stages: it first constructs a physics-informed prior over EOS function space from a pooled ensemble of tabulated hadronic models and then updates the network-weight posterior by placing a multimessenger likelihood directly inside the variational objective. The stellar observations therefore inform the inferred EOS itself.

The physics-informed prior is built from hadronic EOS in the CompOSE database [45]. Rather than regressing pressure directly, the Bayesian network outputs a bounded squared speed of sound, $c_s^2(\epsilon) \in (0, 1)$, and pressure is obtained by integration. Causality, thermodynamic stability, positivity, and monotonicity are therefore structural properties of every core-EOS draw. Three low-density

nuclear anchors and a broad perturbative-QCD anchor set the normalization of the learned functions. Each posterior core EOS is joined to an SLy4 crust and propagated through the exact Tolman-Oppenheimer-Volkoff (TOV-)plus-tidal solver [46–48].

To perform the observational update, we use a differentiable emulator of the EOS-to-observable map inside stochastic variational inference (SVI). The likelihood combines the NICER mass-radius measurements of PSR J0030 + 0451, PSR J0740 + 6620, and PSR J0437-4715 [5–7,11,49], the GW170817 tidal-deformability constraint [4], and the observed $2M_\odot$ maximum-mass bound [9–11]. Gradients through the emulator update the Bayesian neural network (BNN) weights during SVI, while all reported mass-radius and mass-tidal-deformability results are subsequently recomputed with the exact TOV-plus-tidal solver.

Unlike low-dimensional parametric families such as piecewise polytropes and spectral expansions [50,51], the resulting model defines a posterior directly over EOS functions. Relative to our previous Bayesian power-regression model with heteroscedastic errors (BPR-HE) analysis [52], the present framework removes the fixed analytic pressure ansatz, enforces core-EOS admissibility by construction, and incorporates current multimessenger information as an explicit likelihood rather than an a posteriori comparison. Its principal contributions are therefore the structurally admissible speed-of-sound parametrization, the differentiable observational update of the EOS posterior, and the unified propagation of the updated ensemble to both the M - R and M - Λ planes.

A. Equation of state

The 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 nonperturbative regime [1,53].

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 [1,2,53].

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, and 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 EOS in tabulated form with different density grids, compositions, and matching prescriptions between the crust and the core. Community resources such as the CompOSE/COMPOSE database standardize access to large ensembles of tabulated EOS and facilitate systematic comparisons and statistical inference [45].

1. Parametric representations and motivation for a nonparametric approach

EOS inference has traditionally relied on low-dimensional parametric representations: piecewise polytropes [50], spectral expansions of the adiabatic index [51], and empirical Taylor expansions around saturation in coefficients such as K_0 , J , and L [54,55]. These forms are computationally convenient, but each imposes structure that can leak into the inferred physics: segment boundaries produce derivative artifacts in the speed of sound, truncated bases can under-resolve or artificially oscillate around localized features, and near-saturation expansions lose formal control well above ρ_{sat} . In Ref. [52], we introduced a BPR-HE, trained on 65 hadronic EOS from the LIGO LALSuite infrastructure [56], which removed segment-boundary artifacts through a continuous power-law mean relation $P(\epsilon) = a\epsilon^b$ with a density-dependent spread. Its remaining limitation is that the mean trend is still tied to a fixed analytic form, which motivates the present physics-informed Bayesian neural network (PI-BNN) framework as a nonparametric Bayesian successor that learns a posterior *distribution over functions*.

II. BEYOND PARAMETRIZATION: LEARNING DISTRIBUTIONS OVER FUNCTIONS

Traditional EOS parametrizations 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, the PI-BNN defines a nonparametric distribution over dense-matter EOS functions. The network maps energy density to a latent function $g_\theta(\epsilon)$, whose sigmoid gives the bounded sound speed $c_s^2(\epsilon)$; pressure is then obtained by integrating c_s^2 . Bayesian inference over the network weights therefore produces a posterior distribution over physically admissible pressure 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. CompOSE database

To train the BNN and anchor its predictions in established nuclear physics, we construct a comprehensive training dataset sourced from the CompStar Online Supernovae Equations of State (CompOSE) database. CompOSE provides a standardized repository of dense matter EOS models derived from various theoretical frameworks, including relativistic mean-field theories and Skyrme interactions.

The training set comprises 98 cold, beta-equilibrated hadronic tables, selected from CompOSE by composition alone. Hybrid tables and models containing deconfined quark matter are excluded, but no cut is imposed on the maximum TOV mass or on any other macroscopic observable. The ensemble therefore serves as a theory-side prior rather than a set preselected to satisfy astrophysical measurements; those measurements enter only through the observational likelihood described in Sec. IV C. Because the training set contains no hybrid-matter models, posterior draws should not be interpreted as spontaneously reconstructing first-order phase transitions absent from the training distribution.

The tabulated ranges are strongly heterogeneous. All tables begin at essentially vacuum densities in the outer crust ($\epsilon \lesssim 10^{-6}$ MeV/fm³), but their maximum tabulated energy densities range from 51 to 7698 MeV/fm³: 13 of the 98 tables terminate below 200 MeV/fm³ and therefore constrain only the vicinity of the crust-core transition. We truncate the training domain at $\epsilon = 2000$ MeV/fm³, safely above the central energy density of the maximum-mass configuration in our posterior ($\epsilon_c \simeq 1460$ MeV/fm³), so no stable stellar configuration probes densities beyond the cut. Not every table reaches this density: 61 of the 98 tables extend to 2000 MeV/fm³, and only 38 extend to the cutoff of 3000 MeV/fm³. Among the tables tabulated at $\epsilon = 1500$ MeV/fm³, the subset that continues to 3000 MeV/fm³ is systematically softer, with median pressure 588 MeV/fm³, compared with 760 MeV/fm³ for the subset terminating between 1500 and 3000 MeV/fm³. Retaining the sparsely populated tail to 3000 MeV/fm³ would therefore skew the pooled high-density training data toward these softer models without informing any stable stellar configuration. Tables ending earlier contribute no points beyond their tabulated ranges, so model coverage of the pooled dataset decreases with density; accordingly, the high-density posterior is controlled increasingly by the physics anchors rather than by the training data.

This density-dependent coverage also clarifies how the training ensemble enters the BNN: the network is trained on the combined set of (ϵ, P) points from all selected tables, without model-identity labels, rather than on complete EOS curves treated as separate objects. The learned distribution should therefore be interpreted as an ensemble-level prior. At any given density, its predictive spread combines the dispersion among the hadronic tables that extend to that density with the epistemic uncertainty of the network weights. Individual posterior draws are consequently smooth EOS functions supported by the pooled ensemble, not reconstructions of any single microscopic model. We return to the implications of this construction for the low-mass region of the M - R posterior in Sec. IV B.

1. 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, quasiuniform 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 is sensitive to the scale of the input and target variables. Because energy density and pressure span several orders of magnitude, we use the global maximum values $\epsilon_{\max} = \max(\epsilon)$ and $P_{\max} = \max(P)$ to define

$$\tilde{\epsilon} = \frac{\epsilon}{\epsilon_{\max}}, \quad \tilde{P} = \frac{P}{P_{\max}}. \quad (1)$$

The network receives $\tilde{\epsilon}$ as input, but its sigmoid output is the dimensionless physical quantity c_s^2 . The pressure is reconstructed by integrating this output over the corresponding physical energy-density grid and is divided by $P_{\max}$ only when compared with the normalized pressure targets. Consequently, no derivative of a normalized pressure surrogate, and no chain-rule rescaling of such a derivative, enters the present formulation. A numerical offset of 10^{-5} is added to normalized arrays where needed to avoid operations at exact zero. All subsequent equations are written in the physical variables ϵ and P .

B. Physics-informed Bayesian neural network

The surrogate is a fully connected network that maps normalized energy density to the latent function $g_\theta(\epsilon)$. Its hidden layers use the smooth SOFTPLUS activation, which keeps the resulting $c_s^2(\epsilon)$ and integrated pressure smooth. A deterministic network would provide only one extrapolation in the data-sparse inner-core regime; we therefore assign probability distributions to the weights and biases so that uncertainty in the network parameters propagates to uncertainty in the EOS and stellar observables. Independent Gaussian weight priors and a mean-field Gaussian variational posterior are used in the present implementation. The training likelihood compares the integrated pressure $P_\theta(\epsilon_i)$ with each tabulated pressure P_i , while the low-density and perturbative QCD (pQCD) anchors enter as additional log-density terms. The multimessenger likelihood introduced in Sec. IV C subsequently updates this physics-informed prior.

1. BNN architecture

The BNN weights and biases θ are treated as random variables, so uncertainty in the latent function propagates to both $c_s^2(\epsilon)$ and the integrated pressure $P_\theta(\epsilon)$.

Rather than regressing the pressure directly, the network outputs the squared speed of sound as a function of energy density, and the pressure is obtained by integration. Concretely, the network $g_\theta(\epsilon)$ maps energy density to a bounded squared sound speed through a sigmoid output activation,

$$c_s^2(\epsilon) = \text{sigmoid}[g_\theta(\epsilon)] \in (0, 1), \quad (2)$$

and the pressure follows from the thermodynamic identity $c_s^2 = (\partial P / \partial \epsilon)_S$,

$$P(\epsilon) = P_0 + \int_0^\epsilon c_s^2(\epsilon') d\epsilon', \quad (3)$$

where P_0 is the surface pressure fixed by the crust match. This parametrization is the central methodological change relative to direct pressure regression. The sigmoid confines the sound speed to $0 < c_s^2 < 1$, so every core draw is causal and thermodynamically stable; because pressure is the integral of a positive quantity, it is also positive and monotonically increasing. No EOS-level acceptance criterion is therefore applied: every posterior core draw is retained. The integral in Eq. (3) is evaluated by the trapezoidal rule on a fixed energy-density grid, so the map $\theta \mapsto P_\theta(\epsilon)$ remains differentiable for gradient-based variational inference.

We define independent Gaussian priors over all weights and biases,

$$p(\theta) = \prod_i \mathcal{N}(\theta_i | 0, \sigma_w^2), \quad (4)$$

with prior scale $\sigma_w = 0.1$. The hidden layers use the smooth SOFTPLUS activation, which keeps $c_s^2(\epsilon)$ continuously differentiable and therefore yields a smooth pressure profile.

To approximate the intractable true posterior $p(\theta|\mathcal{D})$, we use 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):

$$\mathcal{L}_{\text{ELBO}}(\phi) = \mathbb{E}_{q_\phi(\theta)}[\log p(\mathcal{D}|\theta)] - D_{\text{KL}}(q_\phi(\theta) || p(\theta)). \quad (5)$$

The likelihood is evaluated using the integrated pressure,

$$P_i \sim \mathcal{N}(P_\theta(\epsilon_i), \sigma_i^2), \quad \sigma_i = \sigma_{\text{rel}} P_i + \sigma_{\text{floor}}, \quad (6)$$

with $\sigma_{\text{rel}} = 0.15$ and a small numerical floor. This is the same heteroscedastic likelihood specified in Sec. III A.

A multilayer perceptron (MLP) is sufficient because the learned relation is the one-dimensional map $\epsilon \mapsto c_s^2$. The smooth SOFTPLUS hidden activation and sigmoid output yield a continuously differentiable sound-speed function and a smooth integrated pressure. The mean-field Gaussian variational family provides a practical compromise between posterior flexibility and tractability, although its effect on uncertainty calibration is assessed explicitly in Sec. IV C.

C. Physics-informed constraints

A physical equation of state must respect thermodynamic stability and causality at all densities. As described in Sec. II B 1, these requirements are enforced structurally by the c_s^2 -integral parametrization: the sigmoid confines the sound speed to $0 < c_s^2 < 1$, and integrating this positive quantity makes the pressure positive and monotonic. Consequently, no monotonicity nor causality penalties are used, and no EOS-level acceptance criterion is used; every posterior core draw is retained. The remaining physics input concerns normalization rather than admissibility: the EOS is anchored to nuclear physics at low density and to pQCD at asymptotically high density.

In our probabilistic programming framework (PYRO), these anchors are implemented through PYRO.FACTOR statements, which add custom log-density terms directly to the ELBO during Stochastic Variational Inference, steering the variational posterior toward the physically anchored regions of function space while preserving the flexibility of the neural surrogate.

We use the following anchor terms during SVI:

- (1) *Low-density nuclear anchors.* Rather than a single saturation point, we anchor the low-density EOS at three energy densities, $\epsilon \in \{150, 300, 600\}$ MeV/fm³, with target pressures $\{2.9, 24.8, 124.1\}$ MeV/fm³ and Gaussian widths $\{0.5, 3.0, 15.0\}$ MeV/fm³, respectively. Using three anchors rather than one




FIG. 1. Microphysical predictions of the PI-BNN posterior ensemble. Left: pressure as a function of energy density, where the solid blue curve marks the posterior mean and the shaded blue band denotes the nominal 90% variational 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 $c_s^2 = 1$.

constrains the low-density *shape*, which controls the radius of canonical-mass stars through the pressure near $1-2n_0$.

- (2) *Perturbative QCD anchor.* We place the high-density anchor at $\epsilon_{\text{pQCD}} = 10,000 \text{ MeV/fm}^3$, deep in the pQCD informed regime [57], with a target pressure $P_{\text{pQCD}} = 2300 \text{ MeV/fm}^3$ and a broad Gaussian width $\sigma_{\text{pQCD}} = 1000 \text{ MeV/fm}^3$. The target is set *below* the exact conformal value $P = \epsilon/3 \simeq 3333 \text{ MeV/fm}^3$ because the perturbative pressure at these densities is subconformal, and the broad width encapsulates the residual renormalization-scale uncertainty. We stress that this anchor lies far above the central densities reached by any stable stellar configuration in our posterior, so its role is to guide the high-density extrapolation rather than to constrain the neutron-star-density EOS directly; the c_s^2 plateau it induces above $\epsilon \approx 1500 \text{ MeV/fm}^3$ is therefore a prior-controlled extrapolation, not a data-constrained result.

The anchor target pressures are the medians of the hadronic training ensemble at the three densities, and the saturation value is independently consistent with $P(n_0) \approx Ln_0/3$ for the empirical slope range $L \simeq 40\text{--}70 \text{ MeV}$. The widths $\{0.5, 3.0, 15.0\} \text{ MeV/fm}^3$ are set to approximately half the ensemble standard deviation at each anchor density ($\{1.0, 5.6, 31.5\} \text{ MeV/fm}^3$), so that the anchors constrain the network to the central trend of the models while remaining consistent with the intermodel dispersion. The adequacy of this choice is confirmed *a posteriori* by the agreement of the posterior with the ensemble median at all densities (Fig. 1).

Because causality and monotonicity are guaranteed by the c_s^2 -integral parametrization, the anchor terms are the only physics contributions added to the ELBO, and they serve purely to normalize the EOS at the empirically controlled low-density and asymptotic high-density boundaries. The physics-informed objective is therefore an ELBO augmented only by the anchor log densities,

$$\mathcal{J}_{\text{PI}} = \mathcal{L}_{\text{ELBO}} + W_{\text{low}} \sum_a \log \mathcal{N}(P(\epsilon_a) | P_a, \sigma_a^2) + W_{\text{pQCD}} \log \mathcal{N}(P(\epsilon_{\text{pQCD}}) | P_{\text{pQCD}}, \sigma_{\text{pQCD}}^2), \quad (7)$$

where the sum runs over the three low-density anchors a and $P(\epsilon)$ is the integrated pressure of Eq. (3). This implements Bayesian regularization in function space rather than parameter space and requires no monotonicity or causality terms, since those properties are structural.

1. Posterior ensemble

We draw $N = 500$ posterior core EOS $P_{\theta}^{(k)}(\epsilon)$ from the optimized variational distribution $q_{\theta}^*(\theta)$. Because the c_s^2 -integral parametrization guarantees causality, positivity, and monotonicity for every core draw, no posttraining EOS rejection or acceptance criterion is used. Every draw is stitched to the SLy4 crust and propagated through the TOV-plus-tidal solver, so the reported ensemble is not a selected subset of a larger set of variational samples.

2. Crust-core stitching

The structure of low-mass neutron stars ($M \lesssim 1.0M_{\odot}$) is especially sensitive to the outer and inner crust. Although

the pooled CompOSE tables extend to low density, their crust treatments are heterogeneous; we therefore use the same tabulated SLy4 EOS [48] for every realization and use the BNN draw only for the core segment above the matching interface.

Stitching does not impose physical admissibility on the core, which is already guaranteed by construction; it supplies a common low-density segment and an ordered connection to the BNN core. Let $(\epsilon_{\text{crust,max}}, P_{\text{crust,max}})$ denote the terminal point of the SLy4 crust. For each posterior core EOS $P_{\theta}^{(k)}(\epsilon)$, the matched core segment contains the points satisfying

$$\epsilon > \epsilon_{\text{crust,max}}, \quad P_{\theta}^{(k)}(\epsilon) > P_{\text{crust,max}}. \quad (8)$$

The final stitched tabulation concatenates the crust with this matched core segment,

$$\begin{aligned} \mathcal{T}_{\text{stitched}} = & \mathcal{T}_{\text{SLy4}} \\ & \cup \left\{ (\epsilon, P_{\theta}^{(k)}(\epsilon)) : \epsilon > \epsilon_{\text{crust,max}}, \right. \\ & \left. P_{\theta}^{(k)}(\epsilon) > P_{\text{crust,max}} \right\}, \end{aligned} \quad (9)$$

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

This procedure provides an ordered connection between the SLy4 crust and each posterior core segment. Monotone interpolation over the combined tabulation makes $P(\epsilon)$ continuous and preserves its ordering across the interface; it is not used as a substitute for monotonicity or causality penalties, which are absent from the present formulation.

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.

D. Relativistic structure integration

To map the posterior EOS samples into observable mass-radius and tidal-deformability space, we solve the TOV background equations [46,47] in geometric units ($G = c = 1$):

$$\frac{dP}{dr} = -\frac{(\epsilon + P)(m + 4\pi r^3 P)}{r^2(1 - 2m/r)}, \quad (10)$$

$$\frac{dm}{dr} = 4\pi r^2 \epsilon. \quad (11)$$

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 superluminal 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 $y(r)$, 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

$$r \frac{dy}{dr} + y^2 + yF(r) + r^2 Q(r) = 0, \quad y(0) = 2, \quad (12)$$

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

$$F(r) = \frac{1 - 4\pi r^2 [\epsilon - P]}{1 - 2m/r}, \quad (13)$$

and

$$\begin{aligned} Q(r) = & \frac{4\pi}{1 - 2m/r} \left[5\epsilon + 9P + \frac{\epsilon + P}{c_s^2} - \frac{3}{2\pi r^2} \right] \\ & - 4 \left[\frac{m + 4\pi r^3 P}{r^2(1 - 2m/r)} \right]^2. \end{aligned} \quad (14)$$

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

$$k_2 = \frac{8C^5}{5} (1 - 2C)^2 \frac{2 + 2C(y_R - 1) - y_R}{\Xi(C, y_R)}, \quad (15)$$

$$\begin{aligned} \Xi(C, y_R) = & 2C[6 - 3y_R + 3C(5y_R - 8)] \\ & + 4C^3[13 - 11y_R + C(3y_R - 2)] \\ & + 2C^2(1 + y_R)] \\ & + 3(1 - 2C)^2[2 - y_R \\ & + 2C(y_R - 1)] \ln(1 - 2C). \end{aligned} \quad (16)$$

The corresponding dimensionless tidal deformability is then defined by

$$\Lambda = \frac{2}{3} k_2 C^{-5}. \quad (17)$$

Because the Bayesian posterior produces thousands of EOS realizations, the coupled TOV-plus-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-plus-tidal integration makes the pipeline a joint inference framework for mass, radius, and tidal deformability, rather than an M - R -only mapping.

III. NUMERICAL IMPLEMENTATION

The computational pipeline has two numerical components: probabilistic training of the BNN through SVI [58], 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 [59] and PYRO [60].

Because the BNN parameter space is high dimensional, we perform inference with SVI and approximate the weight posterior with a tractable variational family $q_\phi(\theta)$. The network represents $c_s^2(\epsilon; \theta)$ directly and generates $P_\theta(\epsilon)$ by integration. In practice, PYRO's AUTONORMAL guide provides a mean-field Gaussian approximation over the weights and biases.

The variational parameters ϕ are optimized by maximizing the 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 anchor terms summarized in Sec. II C and the observational likelihood of Sec. IV C are implemented as PYRO.FACTOR contributions to the ELBO. In the present formulation the squared speed of sound is the direct network output rather than a derivative, so no autodifferentiation of the pressure is needed to evaluate thermodynamic admissibility; the pressure entering the anchor and likelihood terms is obtained by trapezoidal integration of $c_s^2(\epsilon)$ on a fixed grid, which keeps the full map differentiable for gradient-based inference.

For each posterior EOS draw, we solve the coupled TOV-plus-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 $\Delta r_{\max} = 0.2$ 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 MLP with one input node, two hidden layers of 64 neurons each, and one output node. The hidden layers use SOFTPLUS, and the output uses a sigmoid to produce c_s^2 directly. The smooth hidden activation prevents artificial structure in the sound-speed function and in the pressure obtained by integration.

All weights and biases are assigned independent Gaussian priors $\mathcal{N}(0, 0.1)$. Variational training is performed with TRACE_ELBO Stochastic Variational Inference using the ADAM optimizer with learning rate 0.01 for 20,000 steps. The likelihood is heteroscedastic: each training point is assigned a Gaussian noise proportional to its pressure, $\sigma_i = \sigma_{\text{rel}} P_i$ with $\sigma_{\text{rel}} = 0.15$ (plus a small floor), so that all densities are constrained at comparable relative precision rather than being dominated by the high-pressure points. The value $\sigma_{\text{rel}} = 0.15$ is somewhat tighter than the raw intermodel dispersion of the hadronic ensemble (≈ 0.25 across the core), a deliberate choice that makes the network follow the central trend of the training models closely while leaving the anchors and the observational likelihood their intended roles.

The anchor locations and target values are those given in Sec. II C: three low-density nuclear anchors at $\epsilon \in \{150, 300, 600\}$ MeV/fm³ with targets $\{2.9, 24.8, 124.1\}$ MeV/fm³ and widths $\{0.5, 3.0, 15.0\}$ MeV/fm³ and a single pQCD anchor at $\epsilon_{\text{pQCD}} = 10,000$ MeV/fm³ with target $P_{\text{pQCD}} = 2300$ MeV/fm³ and width $\sigma_{\text{pQCD}} = 1000$ MeV/fm³. No monotonicity or causality penalties appear, since those properties are enforced structurally by the c_s^2 -integral parametrization (Sec. II B 1); the anchor weights therefore set only how tightly the posterior is held to the empirical low-density shape and the pQCD asymptote. The pressure entering each anchor term is the integrated pressure of Eq. (3), evaluated on a fixed grid of 600 energy-density points spanning from the surface up to the pQCD anchor density by the trapezoidal rule.

IV. RESULTS AND DISCUSSION

A. Microphysical posterior and thermodynamic stability

The primary objective of the PI-BNN is to map the epistemic uncertainty of the dense-matter EOS while enforcing fundamental physical limits. Figure 1 presents the microscopic results of the PI-BNN posterior ensemble.

In the left panel, the reconstructed pressure $P(\epsilon)$ is shown alongside the nominal 90% variational credible interval. At lower densities, the network predictions are tightly constrained by the hadronic training ensemble and the three low-density nuclear anchors. 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 [1,53]. As the energy density increases into the ultradense core regime ($\epsilon \gtrsim 1000$ MeV/fm³), the variational band naturally widens. This broadening reflects increasing epistemic uncertainty in regions with sparse training data, where the inference is guided primarily by pQCD boundary information.

The right panel of Fig. 1 shows the squared speed of sound, $c_s^2 = (\partial P / \partial \epsilon)_S$, which is the direct network output in the present formulation, plotted against the energy density. The mean c_s^2 rises steeply from the crust-core region, crosses the conformal value $c_s^2 = 1/3$ near $\epsilon \approx 500$ MeV/fm³, and levels off at a plateau of $c_s^2 \approx 0.58$ – 0.61 above $\epsilon \approx 1000$ MeV/fm³. The posterior shows a slight, statistically mild decline of the mean above $\epsilon \approx 1500$ MeV/fm³; we note that the parametrization does not forbid such a nonmonotonic c_s^2 , since only the pressure is constrained to be monotonic. By construction c_s^2 remains below the causality limit $c_s^2 = 1.0$ for every posterior draw. Physically, the plateau near $c_s^2 \approx 0.6$ indicates moderately stiff matter, well above the conformal value, as required to support massive stars while remaining causal.

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. Above $\epsilon \approx 500$ MeV/fm³, the speed of sound rises past the conformal value and settles onto a plateau near $c_s^2 \approx 0.6$, corresponding to the additional pressure support needed to sustain stars near the observed $2M_\odot$ threshold. At the highest densities reached in the posterior, the mean c_s^2 shows a mild decline rather than indefinite stiffening, indicating that the inferred EOS balances massive-star support with the high-density guidance from the pQCD anchor 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. Mass-radius and tidal-deformability posterior

By integrating the coupled TOV-plus-tidal equations over the ensemble of causally stable EOS samples, we project the microphysical uncertainties into the observable planes shown in Fig. 2. 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 nominal 90% variational

credible bands are strictly truncated at the median maximum mass of the posterior, preventing statistical artifacts associated with the collapse of softer EOS at lower masses.

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.

The mass-radius band remains well behaved at low masses, with no anomalous tilt toward small radii; the smooth, strictly monotonic c_s^2 -integral EOS yields a regular low-mass branch across the ensemble.

We first describe the physics-only (prior) posterior and then, in Sec. IV C, its observational update; Fig. 2 shows both bands overlaid in each plane.

1. Mass-radius plane

The left panel shows that the physics-only PI-BNN posterior extends above the observed $2.0M_\odot$ lower bound without any ad hoc high-density stiffening. The nominal 90% variational credible interval overlaps with the NICER credible regions for PSR J0030 + 0451, PSR J0740 + 6620, and PSR J0437-4715, which are shown for comparison. The physics-only band gives a canonical radius $R_{1.4} = 13.41$ km and a maximum mass $M_{\text{max}} \gtrsim 2.0M_\odot$.

2. Tidal-deformability plane

The right panel shows that tidal deformability is a central posterior observable. The posterior displays the expected decrease of Λ with increasing mass. The physics-only band gives $\Lambda_{1.4} \simeq 500$; the GW170817 reference value is $\Lambda_{1.4} = 190^{+390}_{-120}$ (90% CI). The prior value reflects the familiar tension between the low GW170817 deformability and the stiffness required to support $2M_\odot$ pulsars; the observational update in Sec. IV C lowers the posterior deformability toward the gravitational-wave constraint.

C. Observational inference of the equation of state

The construction described so far defines a physics-informed prior over EOS functions. We now perform a genuine observational inference by including current multimessenger data directly in the variational objective, so that the network weights—and hence the posterior EOS itself—are updated by the data rather than merely reweighted after the fact.

The obstacle to placing an astrophysical likelihood inside the ELBO is that the map from an EOS to its observables (M, R, Λ) requires integrating the TOV-plus-tidal system, which is not differentiable through the standard adaptive solver. We overcome this with a lightweight differentiable emulator: a small deterministic multilayer perceptron that maps a feature vector of the EOS—the pressures at a fixed set of reference densities—to the mass-radius and mass-tidal-deformability curves and the maximum mass.




FIG. 2. Joint observable-plane predictions from the PI-BNN, showing the physics-only prior (blue) and the observationally updated posterior (salmon/green) as nested nominal 90% variational credible bands, each truncated at the median maximum mass to avoid survivorship bias. Left: the mass-radius (M - R) plane, with the $2.0M_{\odot}$ line and the NICER measurements for PSR J0030 + 0451, PSR J0740 + 6620, and PSR J0437-4715. Right: the mass-tidal-deformability (M - Λ) plane, with the GW170817 constraint. The contraction and shift of the posterior band relative to the prior visualizes the information added by the multimessenger data: the canonical radius moves from $R_{1.4} = 13.41$ to $12.74^{+0.97}_{-0.73}$ km and the tidal deformability from $\Lambda_{1.4} \simeq 500$ to 428^{+249}_{-130} [nominal 90% variational credible interval (CI)], while the maximum mass remains above $2M_{\odot}$.

The emulator is trained once on the physics-only posterior ensemble, augmented with uniformly rescaled variants of each sampled EOS so that it learns the correct, monotonic dependence of the observables on the overall stiffness across the range explored during inference. It reproduces the results of exact solver within 0.15 km at $R_{1.4}$ and is fully differentiable, so the observational log likelihood can be added as a PYRO.FACTOR term evaluated on the emulated observables during SVI.

The observational likelihood combines the NICER mass-radius measurements of PSR J0030 + 0451, PSR J0740 + 6620, and PSR J0437-4715; the GW170817 tidal-deformability constraint at $1.4M_{\odot}$; and a one-sided penalty enforcing the $2M_{\odot}$ maximum-mass bound. To keep the inference stable, we warm-start it from the converged physics-only posterior and ramp the observational weight from 0 to 1 over the first few hundred steps. All observables reported below and summarized in Table I are recomputed with the exact TOV-plus-tidal solver on the resulting posterior ensemble; the emulator is used only to make the training likelihood differentiable.

Table I summarizes the effect of the observational update. The data pull the canonical-mass radius from $R_{1.4} = 13.41$ km in the physics-only prior to $R_{1.4} = 12.74^{+0.97}_{-0.73}$ km (nominal 90% variational CI) in the observationally updated posterior and give $\Lambda_{1.4} = 428^{+249}_{-130}$, while the maximum mass remains above $2M_{\odot}$. The updated posterior remains within the dispersion of the hadronic training ensemble at all

densities, so the softening it represents is data driven rather than an extrapolation beyond the models. As an independent check, importance-reweighting the prior ensemble by the same likelihood yields a consistent canonical radius ($R_{1.4} \simeq 12.9$ km), and a leave-one-out analysis shows that no single measurement dominates the result, the largest individual effect being the removal of PSR J0437-4715 (+0.2 km). The shift between the prior and posterior bands, shown in both observable planes in Fig. 2, is a direct measure of the information content of the current multimessenger dataset. Because the emulator makes the likelihood differentiable and cheap to evaluate, this update requires no re-resolution of the stellar-structure equations during training and can be repeated, or extended with additional likelihoods, as measurements improve.

We assessed the calibration of the variational posterior through a held-out coverage test, retraining on 80% of the hadronic tables and checking how often the resulting bands

TABLE I. Canonical-mass observables before and after the observational update, computed with the exact TOV-plus-tidal solver on the posterior ensemble. Quoted uncertainties are nominal 90% variational credible intervals.

	$R_{1.4}$ (km)	$\Lambda_{1.4}$	$M_{\max}$ ($M_{\odot}$)
Prior (physics only)	13.41	$\simeq 500$	$\gtrsim 2.0$
Posterior (+ obs.)	$12.74^{+0.97}_{-0.73}$	428^{+249}_{-130}	$\gtrsim 2.0$

contain the withheld EOS. The nominal 90% variational interval contains only $\approx 15\%$ of the held-out EOS over the density range constrained by the data, demonstrating substantial underdispersion. The probability-integral transform is centred (mean 0.50), which indicates no detectable systematic rank shift toward softer or stiffer EOS in this test but does not establish that the point estimates are unbiased. We therefore label the reported PI-BNN bands as *nominal* variational credible intervals and interpret their widths as lower bounds on the uncertainty represented by the present mean-field approximation. Richer variational families or sampling-based inference, including Hamiltonian Monte Carlo, may reduce this underdispersion but require additional computational and convergence studies.

V. CONCLUSIONS

In this work, we introduced a physics-informed Bayesian neural network 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 parametrizing the EOS through its speed of sound, $c_s^2(\epsilon) = \text{sigmoid}[g_\theta(\epsilon)]$, and obtaining the pressure by integration, the method guarantees causality, thermodynamic stability, positivity, and monotonicity for every core draw. This eliminates monotonicity and causality penalties as well as EOS-level postsampling rejection: every posterior core draw is retained. Anchor terms at three low-density nuclear points and at the perturbative-QCD boundary fix the normalization of the EOS, and the BNN-to-SLy4 stitching yields stable TOV solutions across both low-mass and high-mass stars.

Beyond constructing this physics-informed prior, we performed a genuine observational inference by embedding the multimessenger likelihood—the NICER mass-radius measurements, the GW170817 tidal-deformability constraint, and the $2M_\odot$ maximum-mass bound—directly in the variational objective. A differentiable emulator of the

TOV-plus-tidal map, validated against the exact solver to within 0.15 km in $R_{1.4}$, makes this update tractable inside the ELBO. The data move the canonical-mass radius from $R_{1.4} = 13.41$ to $12.74^{+0.97}_{-0.73}$ km and the tidal deformability from $\Lambda_{1.4} \simeq 500$ to 428^{+249}_{-130} (nominal 90% variational CI), while the maximum mass remains above $2M_\odot$; the posterior thus overlaps the NICER credible regions, including the refined J0437-4715 measurement, and is consistent with the GW170817 constraint. An independent importance-reweighting analysis and a leave-one-out test confirm that the central result is robust and not driven by any single measurement. The shift between the prior and posterior bands quantifies the information content of the current data, while the held-out coverage test shows that their widths remain underdispersed under the present mean-field approximation. The framework provides a scalable, non-parametric route for incorporating future multimessenger constraints directly into EOS inference, so that improved s-ray, radio, and gravitational-wave measurements can be folded in transparently as they become available.

A limitation of the present analysis is that the training set is restricted to hadronic EOS; explicit phase transitions and hybrid-matter EOS 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. R. V. L. was supported by INCT-FNA (Instituto Nacional de Ciência e Tecnologia, Física Nuclear e Aplicações), research Project No. 464898/2014-5, and also thanks CAPES/FAPERJ/CNPq.

DATA AVAILABILITY

The data that support the findings of this article are openly available in the compose database [61].

-
- [1] M. Oertel, M. Hempel, T. Klähn, and S. Typel, *Rev. Mod. Phys.* **89**, 015007 (2017).
 - [2] O. Komoltsev, R. Somasundaram, T. Gorda, A. Kurkela, J. Margueron, and I. Tews, *Phys. Rev. D* **109**, 094030 (2024).
 - [3] B. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. Lett.* **119**, 161101 (2017).

- [4] B. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. Lett.* **121**, 161101 (2018).
- [5] T. E. Riley *et al.*, *Astrophys. J.* **887**, L21 (2019).
- [6] T. E. Riley *et al.*, *Astrophys. J. Lett.* **918**, L27 (2021).
- [7] D. Choudhury *et al.*, *Astrophys. J. Lett.* **971**, L20 (2024).

- [8] E. Annala, T. Gorda, A. Kurkela, and A. Vuorinen, *Phys. Rev. Lett.* **120**, 172703 (2018).
- [9] J. Antoniadis *et al.*, *Science* **340**, 6131 (2013).
- [10] H. T. Cromartie *et al.* (NANOGrav Collaboration), *Nat. Astron.* **4**, 72 (2019).
- [11] E. Fonseca *et al.*, *Astrophys. J. Lett.* **915**, L12 (2021).
- [12] A. W. Steiner, J. M. Lattimer, and E. F. Brown, *Astrophys. J.* **722**, 33 (2010).
- [13] P. Landry, R. Essick, and K. Chatziioannou, *Phys. Rev. D* **101**, 123007 (2020).
- [14] G. Raaijmakers, S. K. Greif, K. Hebeler, T. Hinderer, S. Nisanke, A. Schwenk, T. E. Riley, A. L. Watts, J. M. Lattimer, and W. C. G. Ho, *Astrophys. J. Lett.* **918**, L29 (2021).
- [15] K. Zhou, L. Wang, L.-G. Pang, and S. Shi, *Prog. Part. Nucl. Phys.* **135**, 104084 (2024).
- [16] G. Papigkiotis, G. Vardakas, A. Likas, and N. Stergioulas, *Phys. Rev. D* **111**, 083056 (2025).
- [17] I. Liodis, E. Smymiotis, and N. Stergioulas, *Phys. Rev. D* **109**, 104008 (2024).
- [18] Y. Fujimoto, K. Fukushima, S. Kamata, and K. Murase, *Phys. Rev. D* **110**, 034035 (2024).
- [19] N. Stergioulas, in *Compact Objects in the Universe*, edited by E. Papantonopoulos and N. Mavromatos (Springer, Cham, 2024), pp. 329–356, [10.1007/978-3-031-55098-0_12](https://doi.org/10.1007/978-3-031-55098-0_12).
- [20] S. Ng, I. Legred, L. Suleiman, P. Landry, L. Traylor, and J. Read, *Classical Quantum Gravity* **42**, 205008 (2025).
- [21] J.-J. Li, Y. Tian, and A. Sedrakian, *Phys. Rev. C* **111**, 055804 (2025).
- [22] X. Dong, H. Shen, J. Hu, and Y. Zhang, *Phys. Rev. D* **112**, 063043 (2025).
- [23] J.-J. Li and A. Sedrakian, *Phys. Rev. C* **112**, 015802 (2025).
- [24] J.-J. Li, Y. Tian, and A. Sedrakian, *Phys. Lett. B* **865**, 139501 (2025).
- [25] G. Papigkiotis, G. Vardakas, and N. Stergioulas, *Phys. Rev. D* **113**, 023018 (2026).
- [26] Y. Fujimoto, K. Fukushima, and K. Murase, *Phys. Rev. D* **98**, 023019 (2018).
- [27] Y. Fujimoto, K. Fukushima, and K. Murase, *Phys. Rev. D* **101**, 054016 (2020).
- [28] Y. Fujimoto, K. Fukushima, and K. Murase, *J. High Energy Phys.* **03** (2021) 273.
- [29] F. Morawski and M. Bejger, *Astron. Astrophys.* **642**, A78 (2020).
- [30] M. Ferreira and C. Providência, *J. Cosmol. Astropart. Phys.* **07** (2021) 011.
- [31] P. G. Krastev, *Galaxies* **10**, 16 (2022).
- [32] M. Ferreira, V. Carvalho, and C. Providência, *Phys. Rev. D* **106**, 103023 (2022).
- [33] P. G. Krastev, *Symmetry* **15**, 1123 (2023).
- [34] A. Thete, K. Banerjee, and T. Malik, *Phys. Rev. D* **108**, 063028 (2023).
- [35] S. Soma, L. Wang, S. Shi, H. Stöcker, and K. Zhou, *J. Cosmol. Astropart. Phys.* **08** (2022) 071.
- [36] S. Soma, L. Wang, S. Shi, H. Stocker, and K. Zhou, *Phys. Rev. D* **107**, 083028 (2023).
- [37] G. c. Gonçalves, M. Ferreira, J. Aveiro, A. Onofre, F. F. Freitas, C. Providência, and J. A. Font, *J. Cosmol. Astropart. Phys.* **12** (2023) 001.
- [38] M.-Z. Han, J.-L. Jiang, S.-P. Tang, and Y.-Z. Fan, *Astrophys. J.* **919**, 11 (2021).
- [39] W. Zhou, J. Hu, Y. Zhang, and H. Shen, *Astrophys. J.* **950**, 186 (2023).
- [40] V. Carvalho, M. Ferreira, T. Malik, and C. Providência, *Phys. Rev. D* **108**, 043031 (2023).
- [41] D. Farrell, P. Baldi, J. Ott, A. Ghosh, A. W. Steiner, A. Kavitar, L. Lindblom, D. Whiteson, and F. Weber, *J. Cosmol. Astropart. Phys.* **02** (2023) 016.
- [42] D. Farrell, P. Baldi, J. Ott, A. Ghosh, A. W. Steiner, A. Kavitar, L. Lindblom, D. Whiteson, and F. Weber, *J. Cosmol. Astropart. Phys.* **12** (2023) 022.
- [43] L. Brandes, C. Modi, A. Ghosh, D. Farrell, L. Lindblom, L. Heinrich, A. W. Steiner, F. Weber, and D. Whiteson, *J. Cosmol. Astropart. Phys.* **09** (2024) 009.
- [44] G. Ventagli and I. D. Saltas, *J. Cosmol. Astropart. Phys.* **01** (2025) 073.
- [45] S. Typel *et al.* (CompOSE Core Team Collaboration), *Eur. Phys. J. A* **58**, 221 (2022).
- [46] R. C. Tolman, *Phys. Rev.* **55**, 364 (1939).
- [47] J. R. Oppenheimer and G. M. Volkoff, *Phys. Rev.* **55**, 374 (1939).
- [48] F. Douchin and P. Haensel, *Astron. Astrophys.* **380**, 151 (2001).
- [49] D. J. Reardon, M. Bailes, R. M. Shannon *et al.*, *Astrophys. J. Lett.* **971**, L18 (2024).
- [50] J. S. Read, B. D. Lackey, B. J. Owen, and J. L. Friedman, *Phys. Rev. D* **79**, 124032 (2009).
- [51] L. Lindblom, *Phys. Rev. D* **82**, 103011 (2010).
- [52] E. V. Chmianski, R. V. Lobato, A. R. Goncalves, and C. A. Bertulani, *Particles* **6**, 198 (2023).
- [53] K. Chatziioannou, H. T. Cromartie, S. Gandolfi, I. Tews, D. Radice, A. W. Steiner, and A. L. Watts, *Rev. Mod. Phys.* **97**, 045007 (2025).
- [54] J. M. Lattimer and Y. Lim, *Astrophys. J.* **771**, 51 (2013).
- [55] J. Margueron, R. Hoffmann Casali, and F. Gulminelli, *Phys. Rev. C* **97**, 025805 (2018).
- [56] LIGO Scientific Collaboration, LIGO algorithm library—lalsuite (2018).
- [57] O. Komoltsev and A. Kurkela, *Phys. Rev. Lett.* **128**, 202701 (2022).
- [58] M. D. Hoffman, D. M. Blei, C. Wang, and J. Paisley, *J. Mach. Learn. Res.* **14**, 1303 (2013), <https://jmlr.org/papers/v14/hoffman13a.html>.
- [59] A. Paszke *et al.*, [arXiv:1912.01703](https://arxiv.org/abs/1912.01703).
- [60] E. Bingham, J. P. Chen, M. Jankowiak, J. M. Obermeyer, N. Pradhan, T. Karaletsos, R. Singh, P. A. Szerlip, P. Horsfall, and N. D. Goodman, *J. Mach. Learn. Res.* **20**, 1 (2019), <https://jmlr.org/papers/v20/18-403.html>.
- [61] CompOSE Core Team *et al.*, CompOSE reference manual: Version 3.01, CompStar Online Supernovae Equations of State, Harmonising the concert of nuclear physics and astrophysics, <https://compose.obspm.fr>.