

# Neural Simulation-Based Inference of the Neutron Star Equation of State directly from Telescope Spectra

Len Brandes,<sup>1,\*</sup> Chirag Modi,<sup>2,3</sup> Aishik Ghosh,<sup>4,5</sup> Delaney Farrell,<sup>6</sup> Lee Lindblom,<sup>7</sup>  
Lukas Heinrich,<sup>1</sup> Andrew W. Steiner,<sup>8,9</sup> Fridolin Weber,<sup>6,7</sup> and Daniel Whiteson<sup>4</sup>

<sup>1</sup>*Technical University of Munich, TUM School of Natural Sciences, Physics Department, 85747 Garching, Germany*

<sup>2</sup>*Center for Computational Astrophysics, Flatiron Institute, New York, NY, 11226, USA*

<sup>3</sup>*Center for Computational Mathematics, Flatiron Institute, New York, NY, 11226, USA*

<sup>4</sup>*Department of Physics and Astronomy, University of California, Irvine, CA 92697, USA*

<sup>5</sup>*Physics Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA*

<sup>6</sup>*Department of Physics, San Diego State University, San Diego, CA 92115, United States*

<sup>7</sup>*Department of Physics, University of California at San Diego, La Jolla, CA 92093, United States*

<sup>8</sup>*Department of Physics and Astronomy, University of Tennessee, Knoxville, Tennessee 37996, USA*

<sup>9</sup>*Physics Division, Oak Ridge National Laboratory, Oak Ridge, Tennessee 37831, USA*

Neutron stars provide a unique opportunity to study strongly interacting matter under extreme density conditions. The intricacies of matter inside neutron stars and their equation of state are not directly visible, but determine bulk properties, such as mass and radius, which affect the star's thermal X-ray emissions. However, the telescope spectra of these emissions are also affected by the stellar distance, hydrogen column, and effective surface temperature, which are not always well-constrained. Uncertainties on these nuisance parameters must be accounted for when making a robust estimation of the equation of state. In this study, we develop a novel methodology that, for the first time, can infer the full posterior distribution of both the equation of state and nuisance parameters directly from telescope observations. This method relies on the use of neural likelihood estimation, in which normalizing flows use samples of simulated telescope data to learn the likelihood of the neutron star spectra as a function of these parameters, coupled with Hamiltonian Monte Carlo methods to efficiently sample from the corresponding posterior distribution. Our approach surpasses the accuracy of previous methods, improves the interpretability of the results by providing access to the full posterior distribution, and naturally scales to a growing number of neutron star observations expected in the coming years.

## CONTENTS

I. Introduction	2
II. Simulated neutron stars	3
A. Equation of state	3
B. Modeling X-ray spectra with XSPEC	3
C. Nuisance parameters	3
III. Previous work	4
IV. Bayesian inference with neural likelihood estimation	5
A. Neural likelihood estimation	5
B. Posterior sampling with Hamiltonian Monte Carlo	5
C. Scaling to multiple observations	6
V. Results	6
A. Example posterior distribution	6
B. Increasing the number of observations	8
C. Average performance on test set	9
VI. Discussion & outlook	10
VII. Conclusion	12

---

\* len.brandes@tum.de

Acknowledgements	13
A. Relativistic mean-field model	13
B. Normalizing flows and simulation-based inference	14
C. Hamiltonian Monte Carlo	15
1. Metropolis-Hastings	15
2. Classical Hamiltonian Monte Carlo	16
3. Implementation	17
References	17

## I. INTRODUCTION

The study of neutron stars provides a window into the fundamental nature of matter under extreme conditions, a long-standing area of research in nuclear and astrophysics [1]. The core of a neutron star contains matter at densities that far surpass those encountered in terrestrial experiments. Under such conditions, matter could exist in various exotic states [2], such as baryons in the form of hyperons and  $\Delta$  isobars [3–7], deconfined quarks [8–11], color superconducting phases [12–14], quarkyonic matter [15, 16], or possibly meson condensates [17–21]. Understanding the physics of matter under the extreme conditions within a neutron star hinges crucially on unveiling the equation of state (EoS), which represents its internal composition through the intricate relationship between pressure and energy density [22]. Knowledge of the EoS is also vital in determining the macroscopic properties of neutron stars, such as their mass and radius, via the relativistic stellar structure equations [23, 24].

Observations of neutron stars have been the primary source for advancing our knowledge of the EoS of superdense matter. In the last two decades, there has been an increasing amount of high-quality data from X-ray and radio emissions, and gravitational waves [25–35]. In particular, spectroscopic measurements from the thermal surface emission of low-mass X-ray binaries in quiescence have provided constraints on the neutron star mass-radius relation. The constraints on the mass-radius relation from observations can be translated into constraints on the EoS of neutron star matter through the use of Bayesian inference techniques [36–50], with a recent interest in complementary methods based on machine learning (ML) [51–72]. Beyond the star’s mass and radius, emitted stellar spectra are also affected by physical quantities like the stellar distance, hydrogen column, and effective surface temperature. Uncertainties on these so-called nuisance parameters must be accounted for when analyzing stellar spectra and making a robust estimation of the equation of state. Traditionally, inference is done in two steps, first by extracting the mass and radius of a neutron star from an observed telescope spectrum, followed by EoS inference from a set of stellar masses and radii [73, 74]. This two-step procedure often requires additional assumptions, though recently it has been demonstrated that a single-step inference is possible while fully propagating uncertainties [71, 72]. While these advancements allow access to the posterior of the EoS parameters by marginalizing nuisance parameters, the computational expense involved has hindered access to the full posterior.

This paper introduces a novel method that provides access to the full posterior distribution of the equation of state and the nuisance parameters directly from telescope spectra. We employ a recently developed simulation-based inference technique known as neural likelihood estimation (NLE) [75] in which normalizing flows [76, 77] use samples of simulated X-ray telescope spectra to learn the likelihood of an observation as a function of the EoS and nuisance parameters. The direct inference of the neutron star EoS from telescope spectra via neural likelihood estimation in contrast to the traditional two-step inference is shown schematically in Fig. 1.

Once trained, the likelihood evaluation is computationally inexpensive, allowing for the calculation of the posterior across the high-dimensional parameter space for inference, marginalization, profiling, or visualization. Since neural networks are, by design, differentiable functions, we can efficiently sample the posterior using state-of-the-art Markov Chain Monte Carlo (MCMC) methods, such as Hamiltonian Monte Carlo (HMC). On simulated test spectra, this method outperforms previous approaches in terms of accuracy and precision of EoS estimation, while improving the interpretability of results with available posterior nuisance parameters. This approach also naturally scales to a growing number of neutron stars, as it does not require retraining to apply to larger datasets.

This paper is organized as follows: Section II reviews the connection between the equation of state of neutron stars and their observed X-ray spectra. This section also outlines the simulation process employed for generating samples of simulated telescope spectra. Sec. III summarizes related previous work to infer the neutron star EoS. Sec. IV then describes our new simulation-based neural likelihood estimation (NLE) approach developed in this paper. The performance of this method is quantitatively compared to previous approaches in Sec. V, with a qualitative discussion on its merits in Sec. VI. Finally, we present our conclusions in Sec. VII.

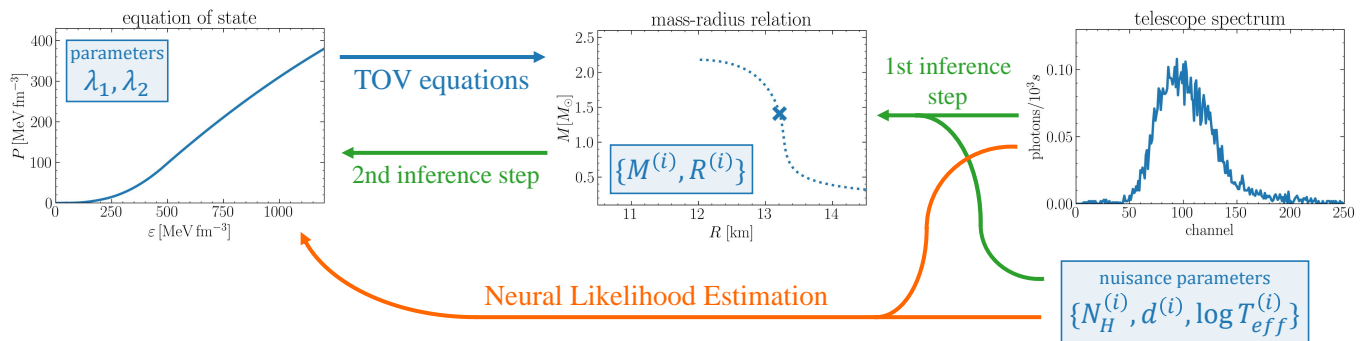


FIG. 1: Traditional inference of EoS parameters (left) from telescope spectra (right) is done by first inferring intermediate mass-radius constraints (green arrows), involving additional implicit assumptions. In contrast, neural likelihood estimation allows for inference of the EoS directly from telescope spectra (orange arrow), robustly accounting for uncertainties.

## II. SIMULATED NEUTRON STARS

Samples of simulated neutron stars with varying values of EoS and nuisance parameters are prepared for training the normalizing flows and evaluation of their performance. Samples from Refs. [71, 72] are used to facilitate direct comparison to previous methods.

### A. Equation of state

Matter inside neutron stars can be described in terms of the EoS, which provides the thermodynamic relationship between pressure  $P$  and energy density  $\varepsilon$  within the star. To generate samples for the EoS, we begin with the relativistic mean field model GM1L [78] (see Appendix A for more details), represented by a second-order spectral expansion as described in Refs. [79, 80]. Thus the full EoS can be described succinctly by the expansion coefficients, which we refer to as  $\lambda_1$  and  $\lambda_2$ . We draw EoS samples by uniformly sampling the coefficients in the intervals  $\lambda_1 \in [4.75, 5.25]$  and  $\lambda_2 \in [-2.05, -1.85]$ .

The underlying physics captured in the star's EoS directly influences bulk properties like mass and radius through the relativistic stellar structure equations. We determine mass-radius relations for each EoS using the Tolman-Oppenheimer-Volkoff (TOV) equations, which determine the structure of non-rotating, spherically symmetric neutron stars in equilibrium [23, 24]. Using the EoS variations described above, the TOV equations are numerically solved to produce mass-radius relations which are sampled to generate the mass  $M$  and radius  $R$  for  $\sim 100$  stars per EoS sample.

### B. Modeling X-ray spectra with XSPEC

Traditional statistical methods infer macroscopic neutron star properties from the emitted X-ray spectra of quiescent low-mass X-ray binaries (LMXBs) by fitting the observed spectrum to well-motivated theoretical models [81–83]. The open-source software XSPEC contains many such models and is widely regarded as the state-of-the-art for spectral fitting [84] and can additionally be used in the generation of simulated spectra. The spectra used in this study were generated using the NSATMOS model in XSPEC, a hydrogen atmospheric model with electron conduction and self-irradiation [83]. Beyond stellar mass and radius, this model also depends on three additional nuisance parameters, described in the next section. The simulated spectra are subjected to the Chandra telescope response function [83] and to Poisson noise corresponding to an observation time of 100 ks.

### C. Nuisance parameters

Each simulated X-ray spectrum from the NSATMOS model depends on the stellar mass, radius, and three additional nuisance parameters: the effective temperature of the surface,  $T_{\text{eff}}$ , the distance to the star,  $d$ , and the hydrogen

column,  $N_H$ , which parameterizes the reddening of the spectrum by the interstellar medium. Values for the nuisance parameters are sampled from ranges with distributions motivated by observation as described in Refs. [25, 26]; details of each range are given in Tab. I.

TABLE I: Distributions and ranges of the equation of state parameters ( $\lambda_1$  and  $\lambda_2$ ) and nuisance parameters ( $N_H$ ,  $d$ , and  $\log T_{\text{eff}}$ ) used to generate the samples of simulated neutron stars. See text for parameter definitions.

Parameter		Distribution	Range
EoS	$\lambda_1$	uniform	[4.75, 5.25]
	$\lambda_2$	uniform	[-2.05, -1.85]
nuisance	$d$	uniform	[2.3, 12.3] kpc
	$N_H$	log uniform	[0.01, 3.16] $10^{21} \text{ cm}^{-2}$
	$T_{\text{eff}}$	exponential	[1, 2] $10^6 \text{ K}$ <sup>a</sup>

<sup>a</sup> This corresponds to a uniform distribution of  $\log(T_{\text{eff}})$  in the range [6, 6.3].

The values of nuisance parameters can be informed by independent observations, which can vary significantly from star to star and provide prior information. We consider three scenarios for nuisance parameter priors, referred to as ‘true’, ‘tight’, and ‘loose’ [71, 72]. In the ‘true’ scenario, the nuisance parameters are known exactly, while the ‘tight’ and ‘loose’ scenarios have narrow or wide Gaussian priors, respectively; see Tab. II for a description of the prior widths for the three scenarios.

TABLE II: Prior distributions on nuisance parameters under three scenarios, ‘true’, ‘tight’, and ‘loose’. Shown are the widths of the Gaussian priors.

Parameter	true	tight	loose
$d$	exact	5%	20%
$N_H$	exact	30%	50%
$\log(T_{\text{eff}})$	exact	$\pm 0.1$	$\pm 0.2$

### III. PREVIOUS WORK

The likelihood,  $p(s|\lambda)$ , of the spectra  $s$  given the EoS parameters  $\lambda$  is not analytically known. Many traditional methods that infer the EoS from telescope spectra approximate the unknown likelihood in a two-step method [73, 74], first inferring posterior distributions  $p(M, R|s)$  for the stellar mass and radius from an observed spectrum. Uncertainties on the nuisance parameters can produce non-trivial [71], occasionally multimodal [26] contours in the posterior probabilities of stellar mass and radius. Subsequently, these posterior distributions, approximated with a Kernel Density Estimator (KDE)<sup>1</sup>, are used as likelihoods in a second step to infer the EoS parameters [37–50]. This is valid only if the  $(M, R)$  priors used in the inference are sufficiently flat [73, 86].

Some previous machine learning (ML) approaches infer the EoS by focusing only on the second of the two steps, starting directly from the posterior probabilities of mass and radius. In Refs. [51–53], neural networks are used to perform regression of EoS parameters from a fixed number of stars described by their masses and radii, where the posterior probabilities are simplified as uncorrelated Gaussians. These neural networks need to be retrained when the number of measurements increases. Several other machine learning methods follow a similar approach with varying architectures [54–69]. The mass and radius standard deviations can be estimated directly from XSPEC by varying the nuisance parameters [71]; a neural network approach that regresses EoS parameters from the resulting mass, radius and uncertainties is referred to below as ‘NN( $M, R$  via XSPEC)’.

However, deep neural networks are capable of analyzing high-dimensional inputs, allowing instead for regression of EoS parameters directly from a set of stellar spectra [71] and effectively removing the intermediate step in the two-step

<sup>1</sup> KDE corresponds to summing a Gaussian kernel at each posterior sample with a standard deviation equal to a bandwidth parameter  $h$  [85].

approach. This method, referred to below as ‘NN(Spectra)’, uses an uncertainty-aware and permutation-invariant neural network and captures the complex correlations between uncertainties, but only produces a point-estimate rather than the full posterior as a function of the EoS parameters and nuisance parameters. Calculation of the full likelihood is intractable, but Ref. [72] showed that two neural networks can be used to replace the unavailable elements, granting access to the likelihood of the expected spectra given stellar mass, radius, and nuisance parameters,  $p(s|M, R, \nu)$ . While this method, referred to as ‘ML-Likelihood<sub>EoS</sub>’, succeeds in obtaining a posterior in EoS, it is computationally expensive to run. This makes it a prohibitively expensive method if access to the nuisance parameter posteriors is also desired. These three approaches are used for comparison in this work.

#### IV. BAYESIAN INFERENCE WITH NEURAL LIKELIHOOD ESTIMATION

We introduce a single-step approach, which uses neural likelihood estimation to directly learn the likelihood of telescope spectra ( $s$ ) as a function of the EoS parameters ( $\lambda$ ) and the nuisance parameters ( $\nu$ ) for a star,  $p(s|\lambda, \nu)$ . This extends the strategy in Ref. [72], which learned  $p(s|M, R, \nu)$ , but avoids needing to integrate over the  $M$ - $R$  plane to achieve a connection to the EoS parameters, saving significant computational complexity. In addition, we apply a Hamiltonian Monte Carlo (HMC) method to efficiently draw samples from the posterior distribution,  $p(\lambda, \nu|s)$ .

##### A. Neural likelihood estimation

Neural likelihood estimation (NLE) [75] is a type of simulation-based inference [87, 88] technique, successfully used for inference with gravitational waves [89–91], particle physics [92–94] and cosmology [95–99] when the likelihood of the observed data are not analytically available but must be estimated from samples of simulated data.

In this approach a neural density estimator ( $q_\Phi$ , with parameters  $\Phi$ ) approximates the likelihood:

$$q_\Phi(s|\lambda, \nu) \approx p(s|\lambda, \nu). \quad (1)$$

Normalizing flows (NF), powerful neural networks capable of modeling complicated probability distributions even in high dimensions [77], are used as the density estimator. A brief introduction to NFs is provided in Appendix B. Once trained, NFs can easily generate new samples as well as estimate the likelihood of a given sample. Specifically, a Masked Autoregressive Flow (MAF) [100] is used for the density estimator  $q_\Phi$ . To fit the parameters  $\Phi$  such that the learned distribution  $q_\Phi$  approximates the target likelihood distribution  $p$ , we minimize the Kullback-Leibler divergence ( $D_{\text{KL}}$ ) between the two distributions, which is equivalent to maximizing the approximate log-probability of the samples generated from the likelihood distribution:

$$\min_{\Phi} D_{\text{KL}}(p(s|\lambda, \nu) || q_\Phi(s|\lambda, \nu)) = \max_{\Phi} \sum_{s_i \sim p(s|\lambda_i, \nu_i)} \log q_\Phi(s_i|\lambda_i, \nu_i). \quad (2)$$

A derivation of this equivalence is given in Appendix B. With samples  $s_i$  drawn from the likelihood via simulation, as described above,

$$s_i \sim p(s|\lambda_i, \nu_i), \quad (3)$$

the `sbi` package [101] is used to train the MAF as a neural likelihood estimator. To make our analysis robust to stochasticities in training the neural network, we perform a hyperparameter search by training 100 different MAFs with different architectures. After training, we use an ensemble average [102] over the  $N = 5$  best-performing density estimators such that the log-likelihood of any given telescope spectrum  $s_0$  is

$$\log p(s_0|\lambda, \nu) \approx \frac{1}{N} \sum_j \log q_{\Phi_j}(s_0|\lambda, \nu). \quad (4)$$

##### B. Posterior sampling with Hamiltonian Monte Carlo

The posterior distribution can be built from the estimated likelihood and prior distributions:

$$p(\lambda, \nu|s) \propto p(s|\lambda, \nu)p(\lambda)p(\nu). \quad (5)$$

The prior distribution on the nuisance parameters,  $p(\nu)$ , is given in Tab. II, while the prior on the EoS parameters,  $p(\lambda)$ , is taken to be uniform in the intervals  $\lambda_1 \in [4.75, 5.25]$  and  $\lambda_2 \in [-2.05, -1.85]$ .

To draw samples from the posterior distribution, sampling algorithms like Markov Chain Monte Carlo (MCMC) or nested sampling can be employed. Given that the normalizing flows, priors and, consequently, the full posterior distribution are differentiable, we draw samples from the posterior using Hamiltonian Monte Carlo (HMC) [103, 104] sampling, which can use the gradient information and scales much more efficiently to high dimensional parameter spaces. For a brief introduction to HMC and our implementation details, see Appendix C.

Unlike standard approaches, our methodology allows for the simultaneous inference of EoS parameters  $\lambda$  and nuisance parameters  $\nu$ , which minimizes assumptions made about these parameters and therefore makes the approach more robust. Additionally, any supplemental information on these parameters coming from other observations can be naturally included in the analysis through the prior distribution  $p(\nu)$  without retraining the neural density estimators.

### C. Scaling to multiple observations

The estimation of the per-star likelihood as a function of the EoS parameters,  $p(s|\lambda, \nu)$ , makes the calculation of the joint likelihood for  $J$  stars straightforward:

$$p(s_{1\dots J}|\lambda, \nu_{1\dots J}) = \prod_j p(s_j|\lambda, \nu_j). \quad (6)$$

Each star has a specific set of nuisance parameters,  $\nu_j$ , such that the posterior is

$$p(\lambda, \nu_{1\dots J}|s_{1\dots J}) \propto \left( \prod_j p(s_j|\lambda, \nu_j)p(\nu_j) \right) p(\lambda). \quad (7)$$

Scaling to multiple observations is straightforward for several reasons. Estimation of the likelihood in the EoS parameters rather than  $M$  and  $R$ , means that no additional integration over the  $M$ - $R$  plane is required. In addition, the likelihood itself is estimated, rather than posteriors which cannot be trivially combined [105–108]. Learning the likelihood conditioned on the nuisance parameters for each star rather than implicitly marginalizing over them allows for the joint likelihood to be simply a product of the individual stellar likelihoods.

A consequence of these choices is that we infer not only the equation of state parameters but also the nuisance parameters corresponding to every star. The stellar nuisance parameter priors are independent and can encode prior information for each star separately, for maximum flexibility and robustness. While this increases the dimensionality of the problem, the availability of gradients of the posterior distribution enables the use of powerful algorithms like HMC, ensuring that the inference remains computationally tractable.

## V. RESULTS

Inference of the full posterior distribution of neutron star EoS and nuisance parameters directly from telescope spectra is now feasible. Following Refs. [71, 72], we apply the method to a set of ten simulated neutron stars for a given point in EoS space, each with three independent nuisance parameters defined above. The complete parameter space comprises two parameters of interest and thirty nuisance parameters. Results are shown for a single representative EoS point, we then demonstrate the scaling of the NLE approach to more observations and finally present a comparison with previous approaches by evaluating the average performance over 100 randomly sampled test points in EoS space.

### A. Example posterior distribution

An example of the posterior distributions marginalized to one and two dimensions is shown in the corner plot in Fig. 2. The figure depicts both EoS parameters,  $\lambda_1$  and  $\lambda_2$ , and the nuisance parameters for the first neutron star,  $N_H^{(1)}$ ,  $d^{(1)}$  and  $\log(T_{\text{eff}}^{(1)})$ . While the marginalized posteriors of the nuisance parameters for the other nine stars are also available, they are not shown here. In all three nuisance parameter scenarios from Tab. II, the EoS parameters are strongly correlated, similar to the log-likelihoods computed in Ref. [72]. The marginal posterior distribution of  $\lambda_1$  is relatively tight compared to its prior range, while for the second parameter  $\lambda_2$  it is not as well constrained compared to the parameter's prior range.

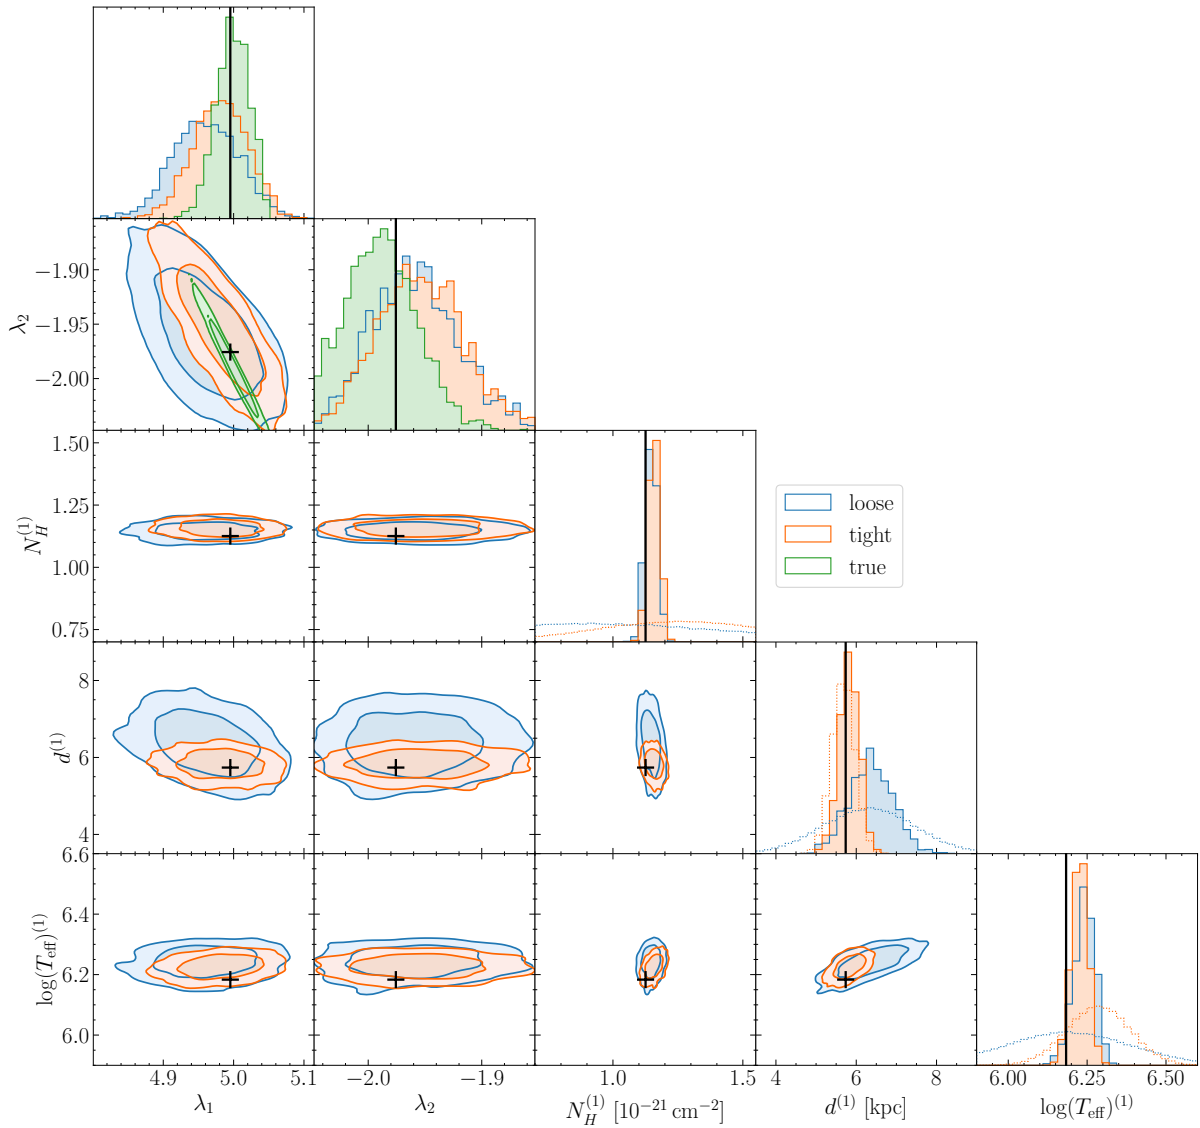


FIG. 2: Corner plot depicting the posterior distribution of the parameters  $\lambda_1$  and  $\lambda_2$  of one example EoS as well as the first 3 (of 30) nuisance parameters  $N_H^{(1)}$ ,  $d^{(1)}$  and  $\log(T_{\text{eff}})^{(1)}$ . The posterior is computed based on the simulated spectra of 10 stars with the nuisance parameters known exactly in the true scenario (green), and known with the uncertainties in Tab. II in the tight (orange) and loose (blue) scenarios. The ground-truth parameter values are depicted as black crosses/lines. The marginal posterior distributions of the nuisance parameters are compared to the respective priors (dotted) of the tight and loose scenarios.

As expected, in the true scenario where the nuisance parameters are exactly known, the marginal posterior distributions are sharply centered around the ground-truth values. In the tight scenario, the uncertainty in the nuisance parameters distributions leads to wider distributions for the EoS parameters. This is further pronounced for the loose case, where less prior information on the nuisance parameters is available. Fig. 2 illustrates that the hydrogen column  $N_H$  as well as the logarithm of the effective surface temperature  $\log(T_{\text{eff}})$  can be significantly constrained from the spectrum data compared to their prior ranges. In the tight scenario, the marginal posterior for the distance  $d$  is almost indistinguishable from the prior, indicating that the telescope spectra do not contribute any more information for this parameter over the tight priors. However, in the loose case, the marginal posterior distribution of  $d$  becomes tighter than the loose prior, which implies that we can indeed extract information about the distance of a neutron star from its X-ray spectrum.

We can transform the posterior distribution for the EoS parameters  $\lambda_1$  and  $\lambda_2$  into 95% (highest density) posterior

credible bands for the pressure as a function of energy density as depicted in Fig. 3. As before, the constraints are much tighter in the true scenario and become increasingly broader in the tight and loose cases. By solving the TOV equations, we can translate the EoS into mass-radius constraints. The 95% posterior credible bands for the radius as a function of mass are depicted in Fig. 3. The credible bands terminate at the 95% upper limits of the maximum mass. Focusing only on the tight case in Fig. 4, there is a very close agreement of the inferred median for  $P(\varepsilon)$  and  $R(M)$  to the ground-truth values.




FIG. 3: Posterior 95% (highest density) credible bands for the pressure as a function of energy density and the radius as a function of mass for the three (true, tight, loose) scenarios. Similar to Fig. 2, the posterior is derived based on the simulated spectra of 10 stars. The ground-truth value for the equation of state and the corresponding mass-radius relation is depicted as a dashed black line. Black dots indicate ground-truth mass-radius values of the 10 simulated stars.




FIG. 4: Median as well as 68% and 95% posterior credible bands for the pressure as a function of energy density and the radius as a function of mass for the *tight* scenario. The ground-truth value for the equation of state and the corresponding mass-radius relation is depicted as a dashed black line. Black dots indicate the mass-radius values of the 10 simulated stars.

## B. Increasing the number of observations

With an anticipated surge in the number of available neutron star observations in the upcoming years, the inference method must be able to scale to a large set of data. In our approach, normalizing flows approximate the likelihood  $p(s|\lambda, \nu)$  per observed star. These likelihoods are then combined to obtain the total likelihood for a set of neutron stars, see Eq. (6). Consequently, the method does not need to assume a fixed number of observed neutron stars, nor a



particular ordering of the stars, and works out-of-the-box for a growing set of observed neutron stars without the need for retraining any networks. To demonstrate this, we present the marginal posterior distributions for one example EoS in Fig. 5, for 5, 10, and 20 observed neutron stars, and with loose uncertainties on the nuisance parameters.




FIG. 5: Corner plot depicting the posterior distribution of the parameters  $\lambda_1$  and  $\lambda_2$  of one example EoS. The posterior is computed based on the simulated spectra of 5 (olive), 10 (blue), or 20 (purple) stars with the nuisance parameters known with the uncertainties in Tab. II of the *loose* scenario. The ground-truth parameter values are depicted as black crosses/lines.

The figure illustrates that the increase of available spectra significantly refines the inference of the EoS parameters. Notably, the transition from 5 to 10 observed spectra has a substantial impact on the posterior constraints, reducing the standard deviation of  $\lambda_1$  by 23% from 0.061 to 0.047, while further increasing the number of measurements to 20 shows a comparatively smaller reduction in the standard deviation by only 6.4% to 0.044 for the given example. For  $\lambda_2$  the increase in accuracy from 10 to 20 observations is even smaller. It is worth noting that in the numerical implementation of Hamiltonian Monte Carlo for posterior sampling, the computation time is predominantly consumed by the evaluation of the likelihood and its gradient. While the availability of more observations increases the per-iteration computational time in sampling the posterior, it also speeds up the convergence of the algorithm. Furthermore, HMC algorithms can easily be scaled to thousands of dimensions, hence we do not anticipate the dimensionality to be a limiting factor in the scaling of our approach.

### C. Average performance on test set

After discussing one example EoS, now we turn to the average performance of NLE with a test set of simulated data from 100 different equations of states. To compare the average performance to the previous ML approaches that infer the neutron star EoS directly from telescope spectra described in Sec. III, we use the same accuracy measure as Refs. [71, 72]. For each EoS in the test set, we simulate 10 spectra with random nuisance parameters. Based on the spectra and the prior nuisance parameter information, we then sample the posterior using the methodology outlined in Sec. IV. From the marginal posterior distributions, similar to the example in Fig. 2, we determine the maximum-a-posteriori estimates (MAP)<sup>2</sup> for the two EoS parameters  $(\lambda_{1,\text{pred}}, \lambda_{2,\text{pred}})$  and compare them to the ground-truth values  $(\lambda_{1,\text{truth}}, \lambda_{2,\text{truth}})$ . The distributions of the differences between the marginal MAP estimate and the ground-truth values,  $(\lambda_{1,\text{pred}} - \lambda_{1,\text{truth}}, \lambda_{2,\text{pred}} - \lambda_{2,\text{truth}})$ , on the test data set are depicted in Fig. 6. As before, the equation

<sup>2</sup> Note that this is the maximum-a-posteriori estimate of the marginal and not of the full posterior.

of state parameters can be maximally constrained in the true nuisance parameter scenario. In all three scenarios, the distributions are centered around zero, indicating that there is no systematic bias in the posterior prediction. Compared to the previous ML analyses in Refs. [71, 72], the distributions of  $\lambda_1$  are much narrower in the tight and loose scenarios (see for example Fig. 8 in [72]).

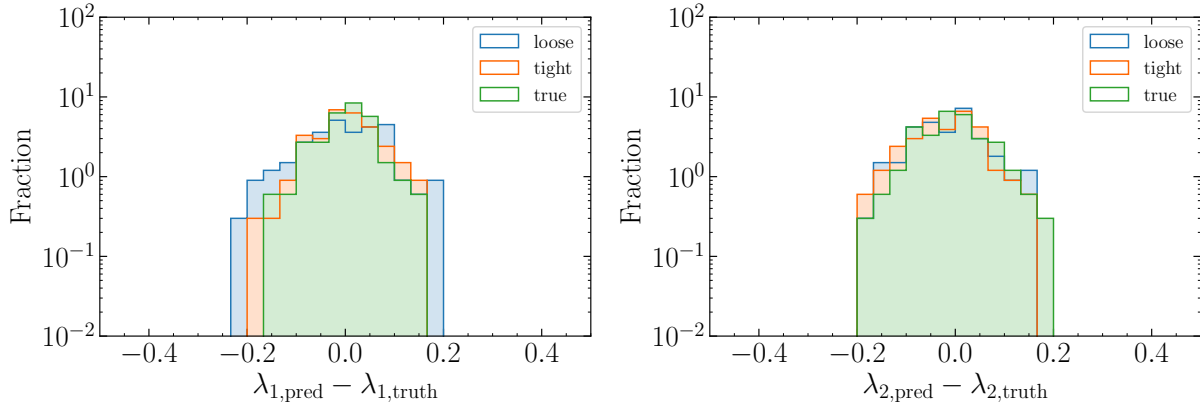


FIG. 6: Distribution of the predicted maximum-a-posteriori estimates versus the ground-truth values for the two equation of state parameters  $\lambda_1$  and  $\lambda_2$ . In the true scenario, the nuisance parameters are fixed to their exact values; in the tight and loose cases, they are drawn from the narrow or wide priors in Tab. II.

To quantitatively compare these distributions to the previous ML analyses, we compute the mean  $\mu$  and standard deviation  $\sigma$  of the distribution of differences in Fig. 6. We can combine both standard deviations into

$$\sigma_{\text{tot}} = \sqrt{\sigma_{\lambda_1}^2 + \sigma_{\lambda_2}^2}. \quad (8)$$

The resulting values are listed in Tab. III and illustrated in Fig. 7. The mean  $\mu$  of the difference between the posterior prediction of  $\lambda_1$  and its ground-truth value is much smaller in the NLE approach compared to previous approaches in all three prior distributions considered in Tab. II. In the true case, the NLE approach performs better than ML-Likelihood<sub>EoS</sub>, from Ref. [72] and NN(Spectra) from Ref. [71] (see Sec. III for more details about these methods). For realistic scenarios of nuisance parameters (loose and tight), NLE outperforms all other methods as quantified by  $\sigma_{\text{tot}}$ , while for the true scenario, it outperforms all methods besides NN( $M, R$  via XSPEC), which uses XSPEC itself for part of the inference and involves simplifying assumptions about the mass-radius uncertainties. Interestingly, from Tab. III it becomes clear that the NLE approach is better than all other approaches to constrain the first EoS parameter  $\lambda_1$ , while for  $\lambda_2$  the accuracy of the XSPEC based two-step approach, with its simplifying assumptions regarding the impact of nuisance parameters, is always much better than all other approaches.

With our neural likelihood estimation approach, we can now, for the first time, infer the posterior for nuisance parameters. The means and standard deviations of the differences between the marginal MAP estimates and the ground-truth values on the test data are listed in Tab. IV. While we obtain excellent constraints on the hydrogen fraction  $N_H$  and the effective surface temperature  $\log(T_{\text{eff}})$ , our ability to constrain the distance  $d$  from the spectra is limited. When going from the tight to the loose scenario, the inference on the distance estimates suffers the most, indicating that these constraints are primarily driven by the prior in the tight case, as seen in Fig. 2. However, this is not as limiting as it might seem because it is much easier to obtain precise constraints on the distance of a neutron star from external measurements than it is for the other nuisance parameters.

## VI. DISCUSSION & OUTLOOK

In addition to the performance gains seen in the preceding sections, the NLE+HMC approach has several advantages relative to previous work.

- (i) The single-step estimation of the likelihood in terms of the EoS avoids collapsing the information into the mass-radius plane as an intermediate step. EoS-dependent quantities like temperature inhomogeneities [109] can impact a star's spectrum but are not captured by the mass and radius. In addition, this avoids assuming the mass-radius posteriors can be used as a likelihood, making simplifying approximations regarding the nature of the intermediate likelihood, or integrating over the  $M$ - $R$  plane.

TABLE III: Average accuracy for the prediction of neutron star EoS parameters  $\lambda_1$  and  $\lambda_2$ . Shown are the means ( $\mu$ ) and standard deviations ( $\sigma$ ) of the distributions in Fig. 6, i.e., of the differences between the predicted maximum-a-posteriori and ground-truth values. Both standard deviations are combined to  $\sigma_{\text{tot}}$  according to Eq. (8). The neural likelihood estimation (NLE) approach is compared to three previous approaches; neural networks that regress the EoS parameters from the spectra (NN(Spectra)) and from  $M, R$  estimates by XSPEC (NN( $M, R$  via XSPEC)), both from Ref. [71], as well as an approach using an approximate likelihood that incorporates two neural networks, ML-Likelihood<sub>EOS</sub>, from [72]. In the true scenario, the nuisance parameters are fixed to their exact values; in the tight and loose cases, they are drawn from the narrow or wide priors in Tab. II.

$p(\nu)$	Method	$\lambda_{1,\text{pred}} - \lambda_{1,\text{truth}}$		$\lambda_{2,\text{pred}} - \lambda_{2,\text{truth}}$		Combined $\sigma_{\text{tot}}$
		$\mu$	$\sigma$	$\mu$	$\sigma$	
true	ML-Likelihood <sub>EOS</sub>	-0.02	0.066	<b>0.01</b>	0.070	0.096
	NN(Spectra)	-0.02	0.066	<b>0.01</b>	0.075	0.099
	NN( $M, R$ via XSPEC)	-0.03	0.065	<b>0.01</b>	<b>0.055</b>	<b>0.085</b>
	<b>NLE</b>	<b>0.00</b>	<b>0.056</b>	<b>-0.01</b>	0.070	0.090
tight	ML-Likelihood <sub>EOS</sub>	-0.02	0.078	0.03	0.081	0.112
	NN(Spectra)	0.02	0.085	-0.02	0.077	0.115
	NN( $M, R$ via XSPEC)	-0.03	0.081	<b>0.01</b>	<b>0.056</b>	0.098
	<b>NLE</b>	<b>0.00</b>	<b>0.066</b>	-0.02	0.071	<b>0.097</b>
loose	ML-Likelihood <sub>EOS</sub>	-0.04	0.089	0.03	0.081	0.120
	NN(Spectra)	-0.03	0.131	<b>-0.01</b>	0.078	0.152
	NN( $M, R$ via XSPEC)	-0.03	0.123	<b>0.01</b>	<b>0.058</b>	0.136
	<b>NLE</b>	<b>0.00</b>	<b>0.085</b>	<b>-0.01</b>	0.074	<b>0.113</b>




FIG. 7: Illustrated mean and standard deviation of the difference between the predicted maximum-a-posteriori values to the ground-truth values for the three different scenarios from Tab. III.

- (ii) Neural likelihood estimation allows for amortization; after training the neural density estimators *once*, the inclusion of additional observations is straightforward, see Sec. V B. In addition, extending to additional stars is inexpensive relative to other methods, which require integrating over estimated mass-radius posteriors to construct likelihoods [39, 50], such as with Kernel Density Estimation techniques.
- (iii) Learning the likelihood instead of the posterior allows combination with likelihoods from other data [110], e.g., constraints from low-energy nuclear theory at small densities [111, 112], perturbative QCD at high densities [113, 114], mass measurements from Shapiro time delays [28–30], mass-radius constraints from analyses of the NICER telescope [31–33] or gravitational wave signals from binary neutron star mergers [34, 35].

TABLE IV: Similar to Tab. III, but with mean and standard deviation of the predicted maximum-a-posteriori estimate versus the ground-truth value for the three nuisance parameters.

Method	$p(\nu)$	$N_{H,\text{pred}} - N_{H,\text{truth}}$		$d_{\text{pred}} - d_{\text{truth}}$		$\log(T_{\text{eff}})_{\text{pred}} - \log(T_{\text{eff}})_{\text{truth}}$	
		$\mu$	$\sigma$	$\mu$	$\sigma$	$\mu$	$\sigma$
NLE	tight	0.00	0.063	-0.06	0.419	0.00	0.034
	loose	0.01	0.070	-0.26	1.149	-0.01	0.047

- (iv) Application of HMC allows for robust exploration of the high-dimensional space of nuisance parameters, which can be of interest in other astrophysical studies.

Improvements to the method could be pursued in several directions. To prevent degradation of performance near the borders of the training data, additional simulated samples could be included beyond the current borders. Another potential improvement may be to additionally condition the inference on the masses  $M$  of the neutron stars, approximating  $p(s|\lambda, \nu, M)$  rather than  $p(s|\lambda, \nu)$ , to further enhance the interpretability and accuracy of the method. In this way, we could infer the neutron star masses and, by solving the TOV equations, radii simultaneously with the EoS. In addition, one could explore alternative parameterizations of the EoS which can describe more complicated phase structures possibly realized inside neutron stars like crossovers of first-order phase transitions, such as piecewise polytropes [115], segment-wise linear interpolations of the speed of sound [116], non-parametric representations based on a Gaussian process [37, 117] or a neural network [118]. In that way, we could constrain the unknown phase structure of strongly interacting matter at large densities [119]. Our approach is conducive to the use of a more complex EoS parameterization.

## VII. CONCLUSION

Neural likelihood estimation (NLE) and Hamiltonian Monte Carlo (HMC) are used in conjunction to allow the first inference of the full posterior for neutron star equation of state and nuisance parameters directly from telescope spectra. In realistic scenarios for nuisance parameter priors, this method outperforms all state-of-the-art methods in terms of prediction accuracy.

Extracting these parameters requires simulation-based inference techniques, as the likelihood is analytically unavailable. Previous methods either relied on inference of the mass and radius as an intermediate step or only provided a point estimate of the EoS parameters. NLE+HMC allows for the full posterior, avoiding simplifying assumptions and information collapse at an intermediate step, and is trivially extendable to multiple stars and other data.

Our study considered three different scenarios for the available prior information on the nuisance parameters coming from additional observations beyond the telescope spectra. With tighter prior constraints, the marginal posterior distributions of the EoS parameters, and accordingly the posterior  $P(\varepsilon)$  and  $R(M)$  credible bands, become much tighter. Both the hydrogen column  $N_H$  and effective surface temperature  $\log(T_{\text{eff}})$  can be tightly constrained from the spectra. For the distance  $d$ , the posterior constraints are driven by the prior information.

These techniques can be extended to NICER or gravitational wave data [120]<sup>3</sup>. Many more gravitational wave events from binary neutron star mergers will be detected in future runs of LIGO, Virgo, and KAGRA. In addition, the gravitational wave database will increase dramatically with the next generation of gravitational wave detectors. For such a large number of expected observations, it is not numerically feasible to compute the likelihood by integrating KDEs [63, 110, 121]. Neural simulation-based inference techniques promise to provide a cost-efficient alternative [90].

This approach can help guide decisions about future observations. Sec. VB demonstrated our method can easily be used to study the impact of additional measurements on the final EoS constraints. It can therefore be used to estimate the relative value of future measurements of one star compared to another, using simulations. Through the assessment of the constraining power of multiple simulated test spectra, we can anticipate which stars provide the most useful additional information required to further constrain the equation of state, thereby guiding decisions for future experimental endeavors.

Based on the results derived in this study on simulated test data, we can conclude that neural likelihood estimation techniques provide a promising avenue for the inference of the neutron star equation of state, not only for X-ray

<sup>3</sup> For the analysis of NICER and gravitational waves, the telescope data is much larger compared to the quiescent LMXBs analyzed here. In that case, it might become necessary to use an additional embedding net to learn summary statistics [101].

measurements of neutron stars but possibly also for gravitational waves. In future studies, this method would be applied to real telescope measurements to draw conclusions about the internal structure of neutron stars.

### ACKNOWLEDGEMENTS

LB thanks Max Dax, Kenji Fukushima, and Koichi Murase for stimulating discussions. LB is supported in part by Deutsche Forschungsgemeinschaft (DFG) and the National Natural Science Foundation of China (NSFC) through funds provided by the Sino-German CRC110 (DFG Grant No. TRR110 and NSFC Grant No. 11621131001), and by the DFG Excellence Cluster ORIGINS. LB thanks the Japan Society for the Promotion of Science (JSPS) and The University of Tokyo for their support and hospitality during his International Fellowship for Research in Japan in the winter and spring of 2024. AG and DW are supported by The Department of Energy Office of Science. FW and DF are supported by the National Science Foundation (USA) under Grant No. PHY-2012152. AWS is supported by NSF AST NSF grant, AST 22-06322 and PHY 21-16686. LL is supported by NSF grant 2012857. LH is supported by the Excellence Cluster ORIGINS, which is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC-2094-390783311.

AG used resources from the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, for this research. We specifically thank Wahid Bhimji and Benjamin Nachman for their support with these computing resources. This collaboration was forged in discussions at the Differentiable and Probabilistic Programming for Fundamental Physics Workshop at Munich and the research was therefore supported by the Munich Institute for Astro-, Particle and BioPhysics (MIAPbP) which is funded by the DFG under Germany's Excellence Strategy – EXC-2094 – 390783311.

### Appendix A: Relativistic mean-field model

In this study, the EoS is based on a relativistic mean-field (RMF) model, which expresses the baryon-baryon interactions through effective meson exchanges. These mesons encompass a scalar meson ( $\sigma$ ) signifying attraction between baryons, a vector meson ( $\omega$ ) representing repulsion, and an isovector meson ( $\rho$ ) essential for elucidating nucleon-nucleon interactions in isospin asymmetric systems. The pion, playing a key role in the long-range description of baryon-baryon interaction, assumes odd parity and consequently disappears in the mean-field approximation. The Lagrangian governing interacting nucleons  $\psi_N$  is expressed as [122]:

$$\mathcal{L}_N = \sum_N \bar{\psi}_N [\gamma_\mu (i\partial^\mu - g_{\omega N}\omega^\mu - \frac{1}{2}g_{\rho N}\boldsymbol{\tau} \cdot \boldsymbol{\rho}^\mu) - (m_N - g_{\sigma N}\sigma)]\psi_N, \quad (\text{A1})$$

where  $g_{\sigma N}$ ,  $g_{\omega N}$ , and  $g_{\rho N}$  represent the meson-nucleon coupling constants.

In the standard RMF approach, meson-nucleon coupling constants are set to reproduce properties of isospin symmetric nuclear matter at nuclear saturation density,  $n_0$ . These properties include the energy per nucleon  $E_0$ , the nuclear incompressibility  $K_0$ , the effective nucleon mass  $m^*/m_N$ , the asymmetry energy  $J$ , the asymmetry energy slope  $L_0$ , and the nucleon potential  $U_N$ . For the GM1L parameter set employed in this study, the respective values are  $n_0 = 0.153 \text{ fm}^{-3}$ ,  $E_0 = -16.3 \text{ MeV}$ ,  $K_0 = 300 \text{ MeV}$ ,  $m^*/m_N = 0.70$ ,  $J = 32.5 \text{ MeV}$ ,  $L_0 = 55.0 \text{ MeV}$ , and  $U_N = -65.5 \text{ MeV}$ . Electrons and muons in neutron star matter are described by the lepton Lagrangian ( $\lambda \in \{e^-, \mu^-\}$ ):

$$\mathcal{L}_L = \sum_\lambda \bar{\psi}_\lambda (i\gamma_\mu \partial^\mu - m_\lambda)\psi_\lambda, \quad (\text{A2})$$

while the meson Lagrangian is given by [122, 123]:

$$\mathcal{L}_M = \frac{1}{2}(\partial_\mu\sigma\partial^\mu\sigma - m_\sigma^2\sigma^2) - \frac{1}{4}\omega_{\mu\nu}\omega^{\mu\nu} + \frac{1}{2}m_\omega^2\omega_\mu\omega^\mu + \frac{1}{2}m_\rho^2\rho_\mu \cdot \rho^\mu - \frac{1}{4}\rho_{\mu\nu} \cdot \rho^{\mu\nu}, \quad (\text{A3})$$

where  $\omega_{\mu\nu} = \partial_\mu\omega_\nu - \partial_\nu\omega_\mu$  and  $\rho_{\mu\nu} = \partial_\mu\rho_\nu - \partial_\nu\rho_\mu$ . To ensure that the RMF model reproduces the empirical values for the nuclear incompressibility and the effective nucleon mass at saturation, additional nonlinear scalar self-interactions need to be included in the Lagrangian [124]:

$$\mathcal{L}_{\text{NL}\sigma} = -\frac{1}{3}b_\sigma m_N [g_{\sigma N}(n)\sigma]^3 - \frac{1}{4}c_\sigma [g_{\sigma N}(n)\sigma]^4, \quad (\text{A4})$$

where  $b_\sigma$  and  $c_\sigma$  are constants determined by the properties of symmetric nuclear matter.

The field equations resulting from the aforementioned Lagrangians must be solved while adhering to the conditions of electrical charge neutrality in neutron star matter:

$$n q_p + \sum_{\lambda} n_{\lambda} q_{\lambda} = 0 , \quad (\text{A5})$$

where  $q_p$  is the electric charge of a proton, and conservation of baryon number:

$$n - \sum_{B=n,p} n_B = 0 . \quad (\text{A6})$$

The meson-field equations combined with the equations for electric charge neutrality and baryon number conservation constitutes a set of five coupled nonlinear equations, to be simultaneously solved to determine the meson mean-fields ( $\bar{\sigma}$ ,  $\bar{\omega}$ ,  $\bar{\rho}$ ) and the neutron and electron Fermi momenta ( $k_n$ ,  $k_e$ ). The proton's Fermi momentum is determined by the condition that neutron star matter is in chemical equilibrium:

$$\mu_p = \mu_n - q_p \mu_e , \quad (\text{A7})$$

where  $\mu_n$  and  $\mu_p$  are the chemical potentials of neutrons and protons respectively, and  $\mu_e$  is the electron chemical potential. The presence of muons in neutron star matter occurs when  $\mu_e \geq \mu_{\mu}$ . For the neutron star matter EoS, we match the GM1L parametrization with the Baym-Pethick-Sutherland (BPS) model [125] for the outer crust and the Baym-Bethe-Pethick (BBP) model [126] for the inner crust.

To limit the number of parameters in the inference procedure, it is advantageous to represent the EoS with only a few parameters. As described in Ref. [71], we accomplish this by constructing a parametric representation based on a spectral fit, where the EoS is represented as a linear combination of basis functions and can therefore be reconstructed using the coefficients of the basis functions. Using a second-order expansion, we construct a spectral fit of the GM1L EoS using the process outlined in Refs. [79, 80]. The two coefficients from the spectral fit are hereafter referred to as  $\lambda_1$  and  $\lambda_2$ . The original coefficients describing the GM1L EoS are then used to generate many EoS scenarios using the expression:

$$\lambda_{\text{generated}} = \lambda_{\text{fit}} \cdot (1 + 2 \cdot c \cdot (\text{ran2} - 0.5)) \quad (\text{A8})$$

where  $\lambda_{\text{generated}}$  represents the newly constructed spectral parameter,  $\lambda_{\text{fit}}$  is the best fit spectral parameter of GM1L, and  $c$  is a scaling parameter set to 0.05.  $\text{ran2}$  are uniformly distributed random numbers in the range 0 to 1 generated by the  $\text{ran2}$  function given in [127]. This process was repeated to create around  $10^3$  different EoS models, each represented by a unique set of spectral coefficients  $\lambda_1$  and  $\lambda_2$ . Following this process, the two parameters are uniformly distributed in the intervals  $\lambda_1 \in [4.75, 5.25]$  and  $\lambda_2 \in [-2.05, -1.85]$ .

## Appendix B: Normalizing flows and simulation-based inference

Normalizing Flows are a class of generative models in machine learning that focus on learning a bijective mapping between a simple base distribution  $\pi(u)$  (usually chosen to be a Gaussian distribution) and a more complex, target distribution  $p(x)$  [77]. The main idea is to model  $p(x)$  by transforming the base distribution  $\pi(u)$  through a series of invertible and differentiable transformations  $f_{\Phi}$  with trainable parameters  $\Phi$ .

Given  $N$  transformations,  $f_{\Phi} = f_{\Phi_1} \circ \dots \circ f_{\Phi_N}$ , we can easily generate samples  $x$  from  $p(x)$  by transforming the samples  $u$  of  $\pi(u)$

$$u \sim \pi(u) , \quad (\text{B1})$$

$$x = f_{\Phi}(u) . \quad (\text{B2})$$

From the change of variables law for probability distributions, we can compute the probability density of the target distribution

$$p(x) = \pi(f_{\Phi}^{-1}(x)) \left| \det \left( \frac{\partial f_{\Phi}^{-1}}{\partial x} \right) \right| . \quad (\text{B3})$$

Consequently, for a fast numerical evaluation of the probability density, the transformations  $f_{\Phi}$  should be chosen such that they are easy to invert and have a Jacobian whose determinant should be fast to compute. One popular choice is Masked Autoregressive Flows (MAF) [100], where each dimension of the data is sequentially transformed conditioned

on the previous dimensions, making the Jacobian of  $f_{\Phi}^{-1}$  triangular by design. MAFs are well suited for our purposes because they are very efficient in computing the probability density  $p(x)$ , but less efficient in generating samples,  $x \sim p(x)$ .

In simulation-based inference (SBI), our goal is to train a conditional normalizing flow to learn the likelihood distribution,  $p(s|\lambda, \nu)$ . This is achieved by minimizing the Kullback-Leibler divergence ( $D_{\text{KL}}$ ), which is a measure of the statistical distance between two probability distributions, between the likelihood  $p(s|\lambda, \nu)$  and the distribution parameterized by the normalizing flow,  $q_{\Phi}(s|\lambda, \nu)$ .  $D_{\text{KL}}$  becomes zero if the distributions are identical. Hence we fit the trainable parameters  $\Phi$  with the following optimization procedure:

$$\min_{\Phi} D_{\text{KL}}(p(s|\lambda, \nu) || q_{\Phi}(s|\lambda, \nu)) = \min_{\Phi} \int ds p(s|\lambda_i, \nu_i) [\log p(s_i|\lambda_i, \nu_i) - \log q_{\Phi}(s_i|\lambda_i, \nu_i)] \quad (\text{B4})$$

$$\approx \min_{\Phi} \sum_{s_i \sim p(s|\lambda_i, \nu_i)} \log p(s_i|\lambda_i, \nu_i) - \log q_{\Phi}(s_i|\lambda_i, \nu_i) \quad (\text{B5})$$

$$= \min_{\Phi} \sum_{s_i \sim p(s|\lambda_i, \nu_i)} -\log q_{\Phi}(s_i|\lambda_i, \nu_i) \quad (\text{B6})$$

$$= \max_{\Phi} \sum_{s_i \sim p(s|\lambda_i, \nu_i)} \log q_{\Phi}(s_i|\lambda_i, \nu_i) , \quad (\text{B7})$$

where the second line is the Monte-Carlo estimator of the integral in the definition of the KL divergence in the first line. The first term in the second line can be dropped as it is constant with respect to the parameters  $\Phi$  of the density estimator. Thus to learn the likelihood distribution, training the density estimator to minimize the KL divergence is equivalent to maximizing the log-probability of the sampled spectra,  $s_i$ , generated from the likelihood distribution via the simulator for given EoS and nuisance parameters,  $(\lambda_i, \nu_i)$ .

### Appendix C: Hamiltonian Monte Carlo

Let  $x \in \mathcal{R}^n$  denote the state (random variable) in a continuous state space. Our goal is to generate samples from a target probability distribution function  $\pi(x)$ . We assume the normalization  $\int dx \pi(x) = 1$ , although all MCMC methods only require un-normalized densities.

#### 1. Metropolis-Hastings

Given the target distribution  $\pi$  and the current state  $x$ , the random-walk Metropolis-Hastings algorithm constructs a Markov chain by sampling a proposal  $y$  from a transition function (proposal distribution)  $t(x, y)$ . The simplest example of such a proposal distribution is a Gaussian distribution centered on  $x$  with some width  $\sigma$ , i.e.,  $y \sim \mathcal{N}(x, \sigma^2)$ . The proposal is then accepted with some probability  $\alpha(x, y)$ , in which case the next state is  $y$ , otherwise it remains  $x$ . We need to identify this acceptance probability to ensure that asymptotically this Markov chain generates samples from the target distribution,  $x, y \sim \pi$ . This is equivalent to ensuring that the target distribution  $\pi$  is invariant under this Markov chain, for example by maintaining detailed balance conditions when accepting the proposed stage

$$\pi(x)t(x, y)\alpha(x, y) = \pi(y)t(y, x)\alpha(y, x) . \quad (\text{C1})$$

Detailed balance means ensuring that the function  $\pi(x)t(x, y)\alpha(x, y)$  is symmetric in exchanging  $x$  and  $y$  ( $x \leftrightarrow y$ ). This equation enforces that the probability of being in the state  $x$ , proposing a transition to state  $y$  and accepting this transition ( $x \rightarrow y$ ) is the same as making the reverse transition ( $y \rightarrow x$ ).

If  $t(x, y)$  is positive everywhere, the above condition is satisfied by the standard Metropolis-Hastings acceptance formula for  $x, y \in S$ ,

$$\alpha(x, y) = \min\left(\frac{\pi(y)t(y, x)}{\pi(x)t(x, y)}, 1\right), \quad (\text{C2})$$

where the denominator is never zero given the above assumptions on  $\pi$  and  $q$ . For each  $x, y \in S$ , either  $\alpha(x, y)$  or  $\alpha(y, x)$  is 1. There are other formulae for  $\alpha$  obeying Eq. (C1), but with lower acceptance rates, meaning they lead to an undesirable higher asymptotic variance for estimated expectations.

## 2. Classical Hamiltonian Monte Carlo

In high dimensions, random walk Metropolis-Hastings as described in the previous section often leads to diffusive behavior and can be extremely inefficient. In this section, we outline the Hamiltonian Monte Carlo (HMC) sampling algorithm which overcomes this by designing more efficient proposal kernels,  $t(x, y)$ , that utilize gradient information [103, 104]. Following standard notation for Hamiltonian dynamics,  $q \in \mathcal{R}^n$  denotes the parameter of interest that is to be sampled. The target density,  $\pi$ , is assumed to be continuous, differentiable, and positive everywhere. To draw samples  $q$  from  $\pi(q)$ , HMC reinterprets the parameters of interest as a position vector with associated potential energy function  $U(q) = -\log \pi(q)$ . We introduce an auxiliary momentum vector  $p \in \mathcal{R}^n$ , which contributes a kinetic energy term  $K(p) = \frac{1}{2}p^T M^{-1}p$ , where  $M$  is some symmetric positive definite mass matrix that we take as fixed. Then the Hamiltonian  $H : \mathcal{R}^{2n} \rightarrow \mathcal{R}$  gives the total energy for the state  $x := (q, p)$ ,

$$H(x) = H(q, p) = U(q) + \frac{1}{2}p^T M^{-1}p . \quad (\text{C3})$$

The state space  $S = \mathcal{R}^{2n}$  is called *phase space*. The dynamical evolution of particles under this Hamiltonian is called Hamiltonian flow and can be simulated by solving Hamiltonian equations.

HMC uses a Markov chain to generate samples  $x$  from the canonical distribution  $\tilde{\pi}$  defined by  $H$ , namely

$$\tilde{\pi}(x) := Z^{-1} e^{-H(x)} = Z^{-1} e^{-U(q)} e^{-\frac{1}{2}p^T M^{-1}p} = Z^{-1} \pi(q) e^{-\frac{1}{2}p^T M^{-1}p} , \quad (\text{C4})$$

where  $Z = \int_{\mathcal{R}^{2n}} dx e^{-H(x)} = (2\pi)^{n/2} \sqrt{\det M}$  is a normalizing constant. Since  $H$  is the sum of potential and kinetic terms, in the Gibbs density  $q$  and  $p$  are independent, with the  $q$ -marginal of  $\tilde{\pi}(x)$  being the target density  $\pi(q)$ . Thus, extracting the first  $n$  coordinates of samples  $x^{(i)}$  from  $\tilde{\pi}$ , one obtains samples from  $\pi$ .

HMC constructs a Markov chain to generate samples from this distribution. Given a current state  $x^{(i)} := (q^{(i)}, p^{(i)})$ , the Markov update comprises two steps:

**Step 1. Gibbs sampling::** Resample the momentum  $p^{(i)}$  from its Gaussian marginal distribution  $p \sim \mathcal{N}(0, M)$ , without changing  $q^{(i)}$ . This randomization step is needed to mix efficiently between different  $H$  values (corresponding to energy level-sets).

**Step 2. Metropolis update::** Given the momentum  $p^{(i)}$ , generate a new Metropolis-Hastings proposal via a deterministic map  $y = F(x)$  which approximates the Hamiltonian flow in Eq. (C3) over a certain time  $T$ , starting from the initial point  $x = (q^{(i)}, p^{(i)})$  and is followed by negation of the final momentum<sup>4</sup>. This map defines the transition kernel  $t$ . This proposal is then accepted with the probability  $\alpha(x, y)$ .

The most commonly used dynamics for the map  $F$  approximating the Hamiltonian flow, i.e., solving the Hamiltonian equations, is the leapfrog (Verlet) integrator. Each leapfrog step, written  $(q', p') = L_\epsilon(q, p)$ , comprises three substeps:

$$\begin{aligned} \bar{p} &\leftarrow p - \frac{\epsilon}{2} \nabla U(q) , \\ q' &\leftarrow q + \epsilon M^{-1} \bar{p} , \\ p' &\leftarrow \bar{p} - \frac{\epsilon}{2} \nabla U(q') . \end{aligned} \quad (\text{C5})$$

where  $\epsilon$  is the step-size. We compose  $n = T/\epsilon$  such steps,  $(L_\epsilon)^n$  to integrate the Hamiltonian flow for time  $T$ . With the momentum-flip operator  $P(q, p) := (q, -p)$ ,  $F = (L_\epsilon)^n P$  (recalling that operators act leftwards) is volume-preserving because each of the three substeps is a shear transformation<sup>5</sup>, and  $F$  is an involution<sup>6</sup> because  $L_\epsilon(q', -p') = (q, -p)$ , which can be verified by reversing the order of the substeps, so  $L_\epsilon P L_\epsilon = P$  and  $((L_\epsilon)^n P)^2 = I$ . In this case, the transition described in Step 2 preserves detailed balance when

$$\alpha(x, y) = \min \left( \frac{\tilde{\pi}(y)}{\tilde{\pi}(x)}, 1 \right) , \quad (\text{C6})$$

so that the distribution  $\tilde{\pi}$  is invariant under Step 2. Since  $\tilde{\pi}$  is also invariant under the Gibbs sampling in Step 1,  $\tilde{\pi}$  is invariant under their composition of both steps, i.e., under each HMC update. Note that failure to approximate well the Hamiltonian flow by the leapfrog integrator *does not* impact detailed balance, although it can drastically reduce the mixing of the Markov chain, and hence the efficiency of the algorithm.

<sup>4</sup> This negation of momentum maintains the invertibility of every step as is necessary for detailed balance, but is not necessary in practice for HMC as it is followed with a Gibbs momentum refresh step.

<sup>5</sup> A *shear* is a map of the form  $(q, p) \mapsto (q + G(p), p)$  or  $(q, p + G(q))$ , and it is easy to check that the determinant of the  $2n \times 2n$  Jacobian derivative matrix is 1.

<sup>6</sup> Involution means  $F^{-1} = F$ , i.e., it is *time reversible*



### 3. Implementation

To draw posterior samples with HMC, we run 16 chains of 2000 samples each. The leapfrog integration in each chain is performed for 40 steps to generate new proposals, with a step-size that is dynamically determined using dual averaging for 300 burn-in steps to achieve an average acceptance probability of 0.65. The mass matrix used in this work is diagonal, except for negative off-diagonal elements between the EoS parameters. In this work, we tuned these values manually based on few initial runs, but in the future this can be automated using variational approximations to the target distribution [128]. We use an importance sampling step to determine initial values for each chain [88].

- 
- [1] G. Baym, T. Hatsuda, T. Kojo, P. D. Powell, Y. Song, and T. Takatsuka, *Rep. Prog. Phys.* **81**, 056902 (2018).
  - [2] D. Blaschke and N. Chamel, Phases of dense matter in compact stars, in *The Physics and Astrophysics of Neutron Stars*, edited by L. Rezzolla, P. Pizzochero, D. I. Jones, N. Rea, and I. Vidaña (Springer International Publishing, 2018) pp. 337–400.
  - [3] L. Tolos and L. Fabbietti, *Prog. Part. Nuc. Phys.* **112**, 103770 (2020).
  - [4] J. J. Li, A. Sedrakian, and F. Weber, *Physics Letters B* **783**, 234 (2018).
  - [5] W. M. Spinella and F. Weber, *Astron. Nachr.* **340**, 145 (2019).
  - [6] G. Malfatti, M. G. Orsaria, I. F. Ranea-Sandoval, G. A. Contrera, and F. Weber, *Phys. Rev. D* **102**, 063008 (2020).
  - [7] A. Sedrakian, J. J. Li, and F. Weber, *Prog. Part. Nucl. Phys.* **131**, 104041 (2023).
  - [8] C. Alcock, E. Farhi, and A. Olinto, *Astrophys. J.* **310**, 261 (1986).
  - [9] J. Madsen, Physics and astrophysics of strange quark matter, in *Hadrons in Dense Matter and Hadrosynthesis*, edited by J. Cleymans, H. B. Geyer, and F. G. Scholtz (Springer Berlin Heidelberg, Berlin, Heidelberg, 1999) pp. 162–203.
  - [10] F. Weber, *Prog. Part. Nucl. Phys.* **54**, 193 (2005).
  - [11] M. Orsaria, H. Rodrigues, F. Weber, and G. A. Contrera, *Phys. Rev. C* **89**, 015806 (2014).
  - [12] M. Alford, *Ann. Rev. Nucl. Part. Sci.* **51**, 131 (2001).
  - [13] M. G. Alford, A. Schmitt, K. Rajagopal, and T. Schäfer, *Rev. Mod. Phys.* **80**, 1455 (2008).
  - [14] J. L. Zdunik and P. Haensel, *Astron. & Astrophys.* **551**, A61 (2013).
  - [15] M. Buballa, *Phys. Rep.* **407**, 205 (2005).
  - [16] L. McLerran and S. Reddy, *Phys. Rev. Lett.* **122**, 122701 (2019).
  - [17] G. Baym, *Phys. Rev. Lett.* **30**, 1340 (1973).
  - [18] D. Kaplan and A. Nelson, *Phys. Lett. B* **175**, 57 (1986).
  - [19] N. K. Glendenning and J. Schaffner-Bielich, *Phys. Rev. C* **60**, 025803 (1999).
  - [20] P. J. Ellis, R. Knorren, and M. Prakash, *Phys. Lett. B* **349**, 11 (1995).
  - [21] A. Ramos, J. Schaffner-Bielich, and J. Wambach, in *Physics of Neutron Star Interiors* (Springer Berlin Heidelberg, 2001) pp. 175–202.
  - [22] T. Kojo, *AAPPS Bull.* **31**, 11 (2021).
  - [23] R. C. Tolman, *Phys. Rev.* **55**, 364 (1939).
  - [24] J. R. Oppenheimer and G. M. Volkoff, *Phys. Rev.* **55**, 374 (1939).
  - [25] J. M. Lattimer and A. W. Steiner, *Astrophys. J.* **784**, 123 (2014).
  - [26] A. W. Steiner, C. O. Heinke, S. Bogdanov, C. K. Li, W. C. G. Ho, A. Bahramian, and S. Han, *Mon. Not. R. Astron. Soc.* **476**, 421 (2018).
  - [27] S. Ascenzi, V. Graber, and N. Rea, arXiv:2401.14930 (2024).
  - [28] J. Antoniadis *et al.*, *Science* **340**, 1233232 (2013).
  - [29] Z. Arzoumanian *et al.*, *Astrophys. J. Suppl. Ser.* **235**, 37 (2018).
  - [30] E. Fonseca *et al.*, *Astrophys. J. Lett.* **915**, L12 (2021).
  - [31] T. E. Riley *et al.*, *Astrophys. J. Lett.* **887**, L21 (2019).
  - [32] M. C. Miller *et al.*, *Astrophys. J. Lett.* **918**, L28 (2021).
  - [33] M. C. Miller *et al.*, *Astrophys. J. Lett.* **887**, L24 (2019).
  - [34] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. X* **9**, 011001 (2019).
  - [35] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Astrophys. J. Lett.* **892**, L3 (2020).
  - [36] G. Raaijmakers *et al.*, *Astrophys. J.* **887**, L22 (2019).
  - [37] I. Legred, K. Chatziioannou, R. Essick, S. Han, and P. Landry, *Phys. Rev. D* **104**, 063003 (2021).
  - [38] L. Brandes, W. Weise, and N. Kaiser, *Phys. Rev. D* **107**, 014011 (2023).
  - [39] L. Brandes, W. Weise, and N. Kaiser, *Phys. Rev. D* **108**, 094014 (2023).
  - [40] S. Huth *et al.*, *Nature (London)* **606**, 276 (2022).
  - [41] S. Altiparmak, C. Ecker, and L. Rezzolla, *Astrophys. J. Lett.* **939**, L34 (2022).
  - [42] Y. Lim and J. W. Holt, *Galaxies* **10**, 99 (2022).
  - [43] M. Marczenko, L. McLerran, K. Redlich, and C. Sasaki, *Phys. Rev. C* **107**, 025802 (2023).
  - [44] E. Annala, T. Gorda, J. Hirvonen, O. Komoltsev, A. Kurkela, J. Nättilä, and A. Vuorinen, *Nature Commun.* **14**, 8451 (2023).

- [45] M.-Z. Han, Y.-J. Huang, S.-P. Tang, and Y.-Z. Fan, *Sci. Bull.* **68**, 913 (2023).
- [46] J.-L. Jiang, C. Ecker, and L. Rezzolla, *Astrophys. J.* **949**, 11 (2023).
- [47] R. Essick, I. Legred, K. Chatziioannou, S. Han, and P. Landry, *Phys. Rev. D* **108**, 043013 (2023).
- [48] D. Mroczek, M. C. Miller, J. Noronha-Hostler, and N. Yunes, arXiv:2309.02345 (2023).
- [49] P. T. H. Pang, L. Sivertsen, R. Somasundaram, T. Dietrich, S. Sen, I. Tews, M. Coughlin, and C. V. D. Broeck, arXiv:2308.15067 (2023).
- [50] H. Koehn, H. Rose, P. T. H. Pang, R. Somasundaram, B. T. Reed, I. Tews, A. Abac, O. Komoltsev, N. Kunert, A. Kurkela, M. W. Coughlin, B. F. Healy, and T. Dietrich, arXiv:2402.04172 (2024).
- [51] Y. Fujimoto, K. Fukushima, and K. Murase, *Phys. Rev. D* **98**, 023019 (2018).
- [52] Y. Fujimoto, K. Fukushima, and K. Murase, *Phys. Rev. D* **101**, 054016 (2020).
- [53] Y. Fujimoto, K. Fukushima, and K. Murase, *J. High Energy Phys.* **2021** (3).
- [54] M. Ferreira and C. Providência, *JCAP* **2021** (07), 011.
- [55] F. Morawski and M. Bejger, *Astron. & Astroph.* **642**, A78 (2020).
- [56] S. Traversi and P. Char, *Astrophys. J.* **905**, 9 (2020).
- [57] P. G. Krastev, *Galaxies* **10**, 16 (2022).
- [58] F. Morawski and M. Bejger, *Phys. Rev. C* **106**, 065802 (2022).
- [59] M. Ferreira, V. Carvalho, and C. Providência, *Phys. Rev. D* **106**, 103023 (2022).
- [60] A. Thete, K. Banerjee, and T. Malik, arXiv:2208.13163 (2022).
- [61] S. Soma, L. Wang, S. Shi, H. Stöcker, and K. Zhou, *JCAP* **2022** (08), 071.
- [62] S. Soma, L. Wang, S. Shi, H. Stöcker, and K. Zhou, *Phys. Rev. D* **107**, 083028 (2023).
- [63] P. G. Krastev, *Symmetry* **15**, 1123 (2023).
- [64] L.-J. Guo, J.-Y. Xiong, Y. Ma, and Y.-L. Ma, arXiv:2309.11227 (2023).
- [65] S. Chatterjee, H. Sudhakaran, and R. Mallick, arXiv:2302.13648 (2023).
- [66] W. Zhou, J. Hu, Y. Zhang, and H. Shen, *Astrophys. J.* **950**, 186 (2023).
- [67] V. Carvalho, M. Ferreira, T. Malik, and C. Providência, *Phys. Rev. D* **108**, 043031 (2023).
- [68] V. Carvalho, M. Ferreira, and C. Providência, arXiv:2401.05770 (2024).
- [69] Y. Fujimoto, K. Fukushima, S. Kamata, and K. Murase, arXiv:2401.12688 (2024).
- [70] K. Zhou, L. Wang, L.-G. Pang, and S. Shi, *Prog. Part. Nucl. Phys.* **135**, 104084 (2024).
- [71] D. Farrell, P. Baldi, J. Ott, A. Ghosh, A. W. Steiner, A. Kavitkar, L. Lindblom, D. Whiteson, and F. Weber, *JCAP* **2023** (02), 016.
- [72] D. Farrell, P. Baldi, J. Ott, A. Ghosh, A. W. Steiner, A. Kavitkar, L. Lindblom, D. Whiteson, and F. Weber, *JCAP* **2023** (12), 022.
- [73] T. E. Riley, G. Raaijmakers, and A. L. Watts, *Mon. Not. R. Astron. Soc.* **478**, 1093 (2018).
- [74] M. C. Miller, C. Chirenti, and F. K. Lamb, *Astrophys. J.* **888**, 12 (2019).
- [75] G. Papamakarios, D. Sterratt, and I. Murray, arXiv:1805.07226 (2019).
- [76] L. Dinh, J. Sohl-Dickstein, and S. Bengio, arXiv:1605.08803 (2016).
- [77] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan, *J. Mach. Learn. Res.* **22** (2021).
- [78] S. Typel, G. Röpke, T. Klähn, D. Blaschke, and H. H. Wolter, *Phys. Rev. C* **81**, 015803 (2010).
- [79] L. Lindblom, *Phys. Rev. D* **82**, 103011 (2010).
- [80] L. Lindblom, *Phys. Rev. D* **97**, 123019 (2018).
- [81] M. Coleman Miller and F. K. Lamb, *Eur. Phys. J. A* **52**, (2016).
- [82] M. C. Miller, *Astrophys. Space Sci. Libr.* **461**, 1 (2020).
- [83] C. O. Heinke, G. B. Rybicki, R. Narayan, and J. E. Grindlay, *Astrophys. J.* **644**, 1090 (2006).
- [84] K. Arnaud, *ASP Conf. Ser.* 101 **17** (1996).
- [85] B. W. Silverman, *Density Estimation for Statistics and Data Analysis* (Chapman and Hall, London, 1986).
- [86] G. Raaijmakers *et al.*, *Astrophys. J. Lett.* **918**, L29 (2021).
- [87] K. Cranmer, J. Brehmer, and G. Louppe, *PNAS* **117**, 30055 (2020).
- [88] J.-M. Lueckmann, J. Boelts, D. S. Greenberg, P. J. Gonçalves, and J. H. Macke, arXiv:2101.04653 (2021).
- [89] A. Delaunoy, A. Wehenkel, T. Hinderer, S. Nissanke, C. Weniger, A. R. Williamson, and G. Louppe, arXiv:2010.12931 (2020).
- [90] M. Dax, S. R. Green, J. Gair, J. H. Macke, A. Buonanno, and B. Schölkopf, *Phys. Rev. Lett.* **127**, 241103 (2021).
- [91] S. R. Green and J. Gair, *MLST* **2**, 03LT01 (2021).
- [92] K. Cranmer, J. Pavez, and G. Louppe, arXiv:1506.02169 (2015).
- [93] J. Brehmer, F. Kling, I. Espejo, and K. Cranmer, *Comput. Softw. Big Sci.* **4**, (2020).
- [94] A. Ghosh, hal-02971995 , 171 (2020).
- [95] J. Alsing, T. Charnock, S. Feeney, and B. Wandelt, *Mon. Not. Roy. Astron. Soc.* , 4440–4458 (2019).
- [96] N. Jeffrey, J. Alsing, and F. Lanusse, *Mon. Not. Roy. Astron. Soc.* **501**, 954 (2020).
- [97] C. Hahn, M. Eickenberg, S. Ho, J. Hou, P. Lemos, E. Massara, C. Modi, A. Moradinezhad Dizgah, B. R.-S. Blacard, and M. M. Abidi, *PNAS* **120**, (2023).
- [98] C. Modi and O. H. E. Philcox, arXiv:2309.10270 (2023).
- [99] C. Modi, S. Pandey, M. Ho, C. Hahn, B. R.-S. Blacard, and B. Wandelt, arXiv:2309.15071 (2023).
- [100] G. Papamakarios, T. Pavlakou, and I. Murray, *NeurIPS* **30** (2017).
- [101] A. Tejero-Cantero, J. Boelts, M. Deistler, J.-M. Lueckmann, C. Durkan, P. Gonçalves, D. Greenberg, and J. Macke, *Journal of Open Source Software* **5**, 2505 (2020).

- [102] J. Hermans, A. Delaunoy, F. Rozet, A. Wehenkel, and G. Louppe, arXiv:2110.06581 (2021).
- [103] R. M. Neal, in Handbook of Markov Chain Monte Carlo (Chapman and Hall/CRC, 2011) Chap. 5.
- [104] M. Betancourt, arXiv:1701.02434 (2017).
- [105] G. Papamakarios and I. Murray, arXiv:1605.06376 (2018).
- [106] D. S. Greenberg, M. Nonnenmacher, and J. H. Macke, arXiv:1905.07488 (2019).
- [107] C. Durkan, I. Murray, and G. Papamakarios, arXiv:2002.03712 (2020).
- [108] J. Hermans, V. Begy, and G. Louppe, arXiv:1903.04057 (2020).
- [109] K. G. Elshamouty, C. O. Heinke, S. M. Morsink, S. Bogdanov, and A. L. Stevens, *Astrophys. J.* **826**, 162 (2016).
- [110] J. Lange, R. O'Shaughnessy, and M. Rizzo, arXiv:1805.10457 (2018).
- [111] C. Drischler, S. Han, and S. Reddy, *Phys. Rev. C* **105**, 035808 (2022).
- [112] J. Keller, K. Hebeler, and A. Schwenk, *Phys. Rev. Lett.* **130**, 072701 (2023).
- [113] T. Gorda, A. Kurkela, R. Paatelainen, S. Säppi, and A. Vuorinen, *Phys. Rev. Lett.* **127**, 162003 (2021).
- [114] O. Komoltsev and A. Kurkela, *Phys. Rev. Lett.* **128**, 202701 (2022).
- [115] S. K. Greif, G. Raaijmakers, K. Hebeler, A. Schwenk, and A. L. Watts, *Mon. Not. R. Astron. Soc.* **485**, 5363 (2019).
- [116] E. Annala, T. Gorda, A. Kurkela, J. Nättilä, and A. Vuorinen, *Nat. Phys.* **16**, 907 (2020).
- [117] P. Landry, R. Essick, and K. Chatziioannou, *Phys. Rev. D* **101**, 123007 (2020).
- [118] M.-Z. Han, J.-L. Jiang, S.-P. Tang, and Y.-Z. Fan, *Astrophys. J.* **919**, 11 (2021).
- [119] L. Brandes and W. Weise, *Symmetry* **16**, 111 (2024).
- [120] G. Gonçalves, M. Ferreira, J. Aveiro, A. Onofre, F. F. Freitas, C. Providência, and J. A. Font, *JCAP* **2023** (12), 001.
- [121] F. Hernandez Vivanco, R. Smith, E. Thrane, P. D. Lasky, C. Talbot, and V. Raymond, *Phys. Rev. D* **100**, 103009 (2019).
- [122] N. Glendenning, Compact Stars: Nuclear Physics, Particle Physics and General Relativity, Astronomy and Astrophysics Library (Springer New York, 2012).
- [123] Y. Gambhir, P. Ring, and A. Thimet, *Ann. Phys.* **198**, 132 (1990).
- [124] J. Boguta and A. Bodmer, *Nucl. Phys. A* **292**, 413 (1977).
- [125] G. Baym, C. Pethick, and P. Sutherland, *Astrophys. J.* **170**, 299 (1971).
- [126] G. Baym, H. A. Bethe, and C. J. Pethick, *Nucl. Phys. A* **175**, 225 (1971).
- [127] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, Numerical Recipes 3rd Edition (Cambridge University Press, 2007) p. 1256.
- [128] C. Modi, C. Margossian, Y. Yao, R. Gower, D. Blei, and L. Saul, arXiv:2307.07849 (2023).