Journal of Cosmology and
Astroparticle Physics

SISSA

PAPER • OPEN ACCESS

Neural simulation-based inference of the neutron star equation of state directly from telescope spectra

To cite this article: Len Brandes *et al* JCAP09(2024)009

View the article online for updates and enhancements.

You may also like

- Distinguishing Binary Neutron Star from Neutron Star–Black Hole Mergers with Gravitational Waves
Hsin-Yu Chen and Katerina Chatziioannou
- Nuclear physics of the outer layers of accreting neutron stars
Zach Meisel, Alex Deibel, Laurens Keek et al.
- Unified neutron star EOSs and neutron star structures in RMF models
Cheng-Jun Xia, Toshiki Maruyama, Ang Li et al.

This content was downloaded from IP address 146.75.146.78 on 05/11/2024 at 18:51

Journal of Cosmology and Astroparticle Physics
An IOP and SISSA journal

Received: March 12, 2024
Revised: June 10, 2024
Accepted: August 10, 2024
Published: September 3, 2024

Neural simulation-based inference of the neutron star equation of state directly from telescope spectra

Len Brandes[a], **Chirag Modi**[b,c], **Aishik Ghosh**[d,e], **Delaney Farrell**[f], **Lee Lindblom**[g], **Lukas Heinrich**[a], **Andrew W. Steiner**[h,i], **Fridolin Weber**[f,g] and **Daniel Whiteson**[d]

[a] *Physics Department, TUM School of Natural Sciences, Technical University of Munich, Garching 85747, Germany*
[b] *Center for Computational Astrophysics, Flatiron Institute, New York, NY 11226, U.S.A.*
[c] *Center for Computational Mathematics, Flatiron Institute, New York, NY 11226, U.S.A.*
[d] *Department of Physics and Astronomy, University of California, Irvine, CA 92697, U.S.A.*
[e] *Physics Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, U.S.A.*
[f] *Department of Physics, San Diego State University, San Diego, CA 92115, U.S.A.*
[g] *Department of Physics, University of California at San Diego, La Jolla, CA 92093, U.S.A.*
[h] *Department of Physics and Astronomy, University of Tennessee, Knoxville, TN 37996, U.S.A.*
[i] *Physics Division, Oak Ridge National Laboratory, Oak Ridge, TN 37831, U.S.A.*

E-mail: len.brandes@tum.de, cmodi@flatironinstitute.org, aishikghosh@cern.ch, dfarrell@sdsu.edu, lindblom@tapir.caltech.edu, lukas.heinrich@cern.ch, awsteiner@utk.edu, fweber@sciences.sdsu.edu, daniel@uci.edu

Abstract: 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

© 2024 The Author(s). Published by IOP Publishing Ltd on behalf of Sissa Medialab. Original content from this work may be used under the terms of the Creative Commons Attribution 4.0 licence. Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

https://doi.org/10.1088/1475-7516/2024/09/009

JCAP09(2024)009

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.

Keywords: Machine learning, neutron stars, Bayesian reasoning, X-ray binaries

ArXiv ePrint: 2403.00287

JCAP09(2024)009

Contents

JCAP09(2024)009

1 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 Δ 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

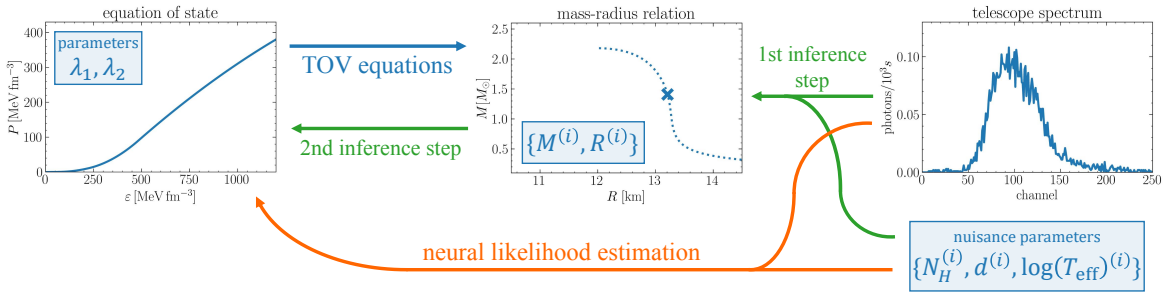

Figure 1. Traditional inference of EoS parameters (left) from telescope spectra (right) is done by first inferring intermediate mass-radius constraints (green arrows), involving additional implicit assumptions. In contrast, neural likelihood estimation allows for inference of the EoS directly from telescope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

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 figure 1.

JCAP09(2024)009

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 2 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. Section 3 summarizes related previous work to infer the neutron star EoS. Section 4 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 section 5, with a qualitative discussion on its merits in section 6. Finally, we present our conclusions in section 7.

2 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.

2.1 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 ε 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 λ_1 and λ_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. For the sampling of stars we choose a log-uniform distribution of central enthalpies [81] for the boundary condition to solve the stellar structure equations. This leads to a higher weighting of larger masses close to the maximum supported mass in the distribution of stars for each EoS.

JCAP09(2024)009

Parameter		Distribution	Range
EoS	λ_1	uniform	$[4.75, 5.25]$
	λ_2	uniform	$[-2.05, -1.85]$
nuisance	d	uniform	$[2.3, 12.3]$ kpc
	N_H	log uniform	$[0.01, 3.16]\, 10^{21}$ cm^{-2}
	$T_{\rm eff}$	exponential	$[1, 2] 10^6$ K,[a]

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

Table 1. Distributions and ranges of the equation of state parameters (λ_1 and λ_2) and nuisance parameters (N_H, d, and $\log(T_{\rm eff})$) used to generate the samples of simulated neutron stars. See text for parameter definitions.

2.2 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 [82–84]. The open-source software XSPEC contains many such models and is widely regarded as the state-of-the-art for spectral fitting [85] 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 [84]. 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 corresponding to the instrument ACIS-S [84, 86] and to Poisson noise corresponding to an observation time of 100 ks.[1]

2.3 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, $T_{\rm eff}$, 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 table 1.

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 table 2 for a description of the prior widths for the three scenarios.

[1] The simulated spectra do not yield detectable photons at high energies. Therefore, we follow previous works [71, 72] to use only the first 250 bins of the telescope spectra.

JCAP09(2024)009

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

Table 2. Prior distributions on nuisance parameters under three scenarios, 'true', 'tight', and 'loose'. Shown are the widths of the Gaussian priors.

3 Previous work

The likelihood, $p(s|\lambda)$, of the spectra s given the EoS parameters λ 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),[2] 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, 88].

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 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 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$_{\rm EOS}$', 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.

[2] KDE corresponds to summing a Gaussian kernel at each posterior sample with a standard deviation equal to a bandwidth parameter h [87].

JCAP09(2024)009

4 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 (λ) and the nuisance parameters (ν) 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, $p(\lambda,\nu|s)$.

4.1 Neural likelihood estimation

Neural likelihood estimation (NLE) [75] is a type of simulation-based inference [89, 90] technique, successfully used for inference with gravitational waves [91–93], particle physics [94–96] and cosmology [97–101] 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_Φ, with parameters Φ) approximates the likelihood:

$$q_\Phi(s|\lambda,\nu) \approx p(s|\lambda,\nu). \tag{4.1}$$

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) [102] is used for the density estimator q_Φ. To fit the parameters Φ such that the learned distribution q_Φ approximates the target likelihood distribution p, we minimize the Kullback-Leibler divergence ($D_{\rm KL}$) between the two distributions, which is equivalent to maximizing the approximate log-probability of the samples generated from the likelihood distribution:

$$\arg\min_\Phi D_{\rm KL}\left(p(s|\lambda,\nu)||q_\Phi(s|\lambda,\nu)\right) = \arg\max_\Phi \sum_{s_i\sim p(s|\lambda_i,\nu_i)} \log q_\Phi(s_i|\lambda_i,\nu_i)\,. \tag{4.2}$$

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

$$s_i \sim p(s|\lambda_i,\nu_i)\,, \tag{4.3}$$

the `sbi` package [103] 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 [104] over the $N = 5$ best-performing density estimators such that the log-likelihood of any given telescope spectrum s_0 is

$$\log p(s_0|\lambda,\nu) \approx \frac{1}{N}\sum_j \log q_{\Phi_j}(s_0|\lambda,\nu)\,. \tag{4.4}$$

JCAP09(2024)009

4.2 Posterior sampling with Hamiltonian Monte Carlo

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

$$p(\lambda, \nu|s) \propto p(s|\lambda, \nu)p(\lambda)p(\nu)\,. \tag{4.5}$$

The prior distribution on the nuisance parameters, $p(\nu)$, is given in table 2, 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 $\lambda_2 \in [-2.05, -1.85]$, the same as the distribution of training samples described in table 1.

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) [105, 106] 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 λ and nuisance parameters ν, 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.

4.3 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:

$$p(s_{1\ldots J}|\lambda, \nu_{1\ldots J}) = \prod_j p(s_j|\lambda, \nu_j)\,. \tag{4.6}$$

Each star has a specific set of nuisance parameters, ν_j, such that the posterior is

$$p(\lambda, \nu_{1\ldots J}|s_{1\ldots J}) \propto \left(\prod_j p(s_j|\lambda, \nu_j)p(\nu_j)\right) p(\lambda)\,. \tag{4.7}$$

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 [107–110]. 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.

JCAP09(2024)009

5 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.

5.1 Example posterior distribution

An example of the posterior distributions marginalized to one and two dimensions is shown in the corner plot in figure 2. The figure depicts both EoS parameters, λ_1 and λ_2, and the nuisance parameters for the first neutron star, $N_H^{(1)}$, $d^{(1)}$ and $\log(T_{\rm eff})^{(1)}$. 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 table 2, the EoS parameters are strongly correlated, similar to the log-likelihoods computed in ref. [72]. The marginal posterior distribution of λ_1 is relatively tight compared to its prior range, while for the second parameter λ_2 it is not as well constrained compared to the parameter's prior range.

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 parameter 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. Figure 2 illustrates that the hydrogen column N_H as well as the logarithm of the effective surface temperature $\log(T_{\rm eff})$ 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 λ_1 and λ_2 into 95% (highest density) posterior credible bands for the pressure as a function of energy density as depicted in figure 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 figure 3. The credible bands terminate at the 95% upper limits of the maximum mass. Focusing only on the tight case in figure 4, there is a very close agreement of the inferred median for $P(\varepsilon)$ and $R(M)$ to the ground-truth values. Note that for this particular example, the mass of one of the simulated stars is very close to the respective maximum supported mass such that a good reconstruction even of the high-density part for this particular EoS is possible. In other cases, the reconstruction might be worse.

JCAP09(2024)009

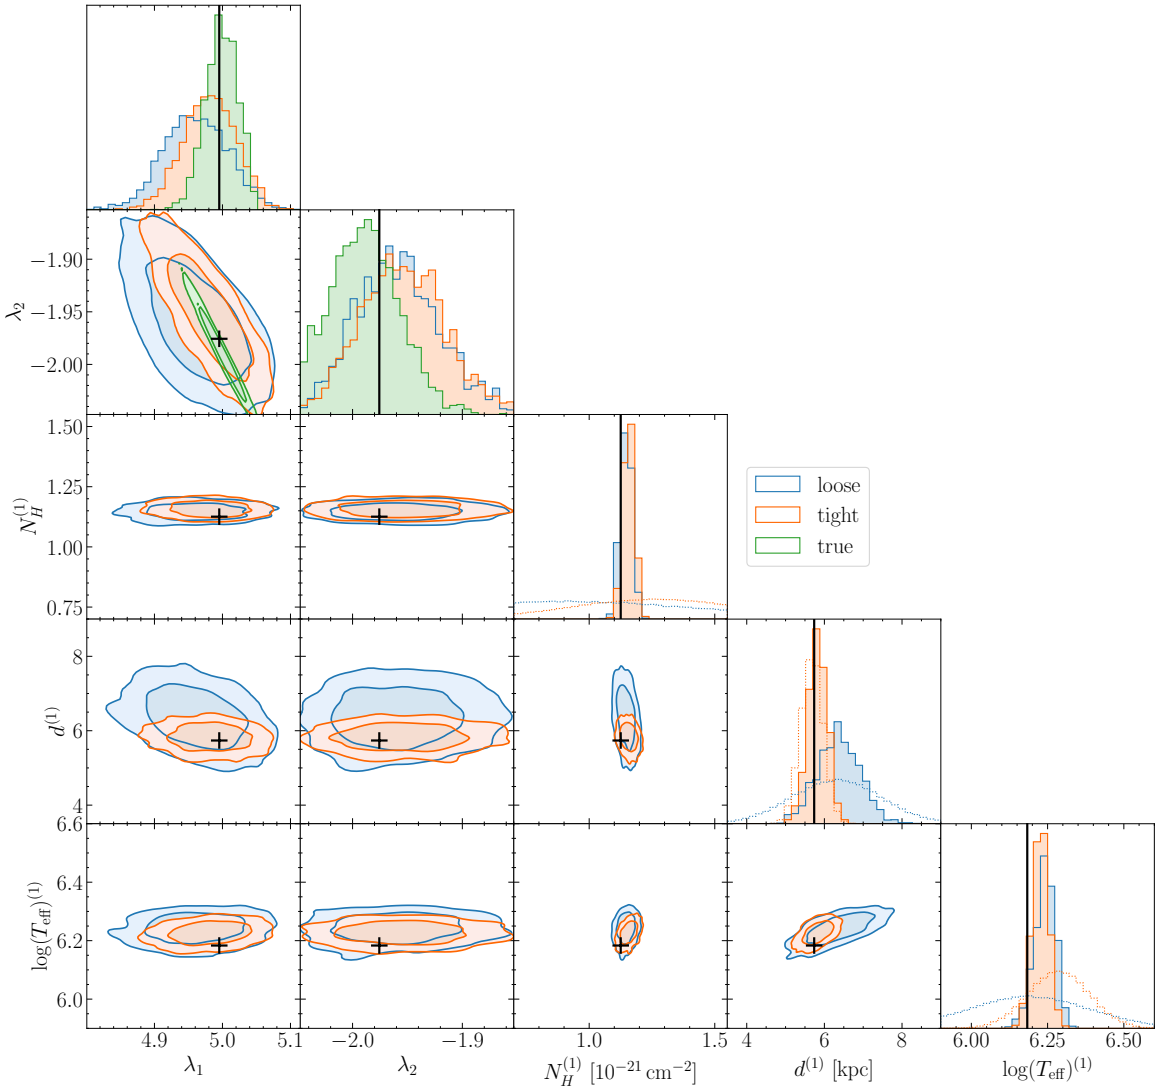

Figure 2. Corner plot depicting the posterior distribution of the parameters λ_1 and λ_2 of one example EoS as well as the first 3 (of 30) nuisance parameters $N_H^{(1)}$, $d^{(1)}$ and $\log(T_{\rm eff})^{(1)}$. The posterior is computed based on the simulated spectra of 10 stars with the nuisance parameters known exactly in the true scenario (green), and known with the uncertainties in table 2 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.

5.2 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. (4.6). 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

JCAP09(2024)009

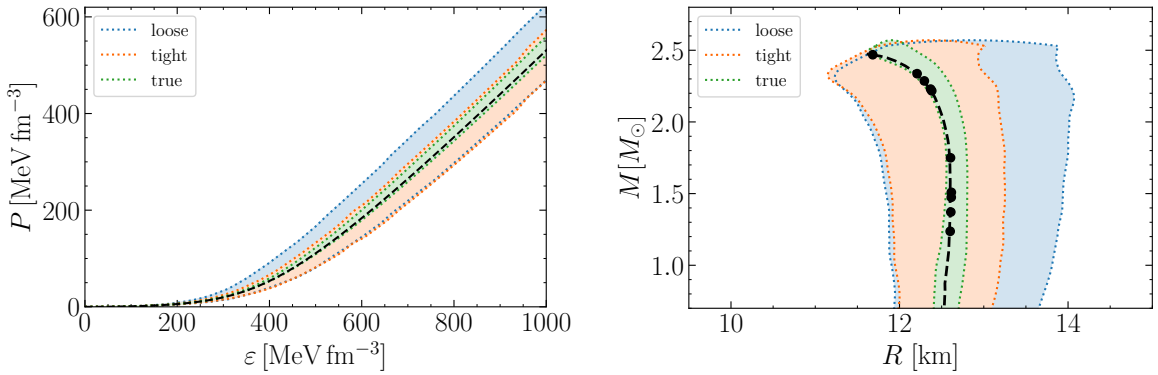

Figure 3. Posterior 95% (highest density) credible bands for the pressure as a function of energy density and the radius as a function of mass for the three (true, tight, loose) scenarios. Similar to figure 2, the posterior is derived based on the simulated spectra of 10 stars. The ground-truth value 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.

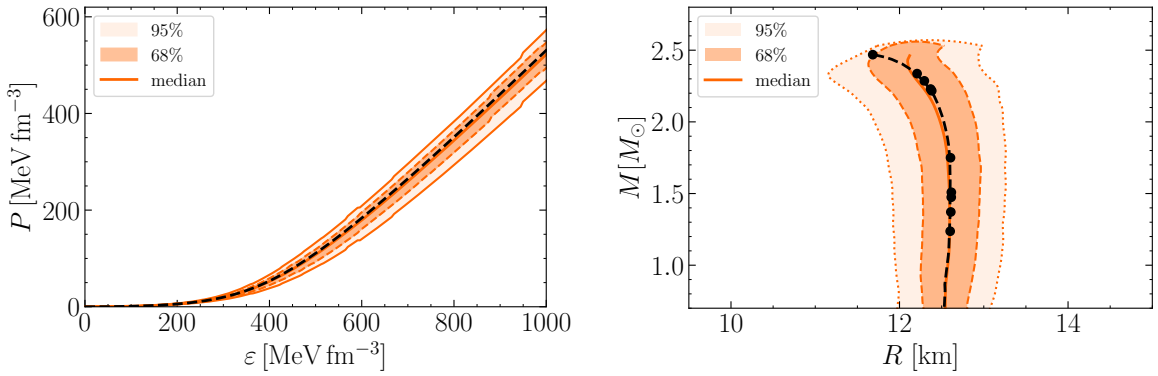

Figure 4. Similar to figure 3, but only for the *tight* scenario the median and the 68% and 95% posterior credible bands are shown for the pressure as a function of energy density and the radius as a function of mass. The ground-truth value for the equation of state and the corresponding mass-radius relation is depicted as a dashed black line. Black dots indicate the mass-radius values of the 10 simulated stars.

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 figure 5, for 5, 10, and 20 observed neutron stars, and with loose uncertainties on the nuisance parameters.

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 λ_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 λ_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

JCAP09(2024)009

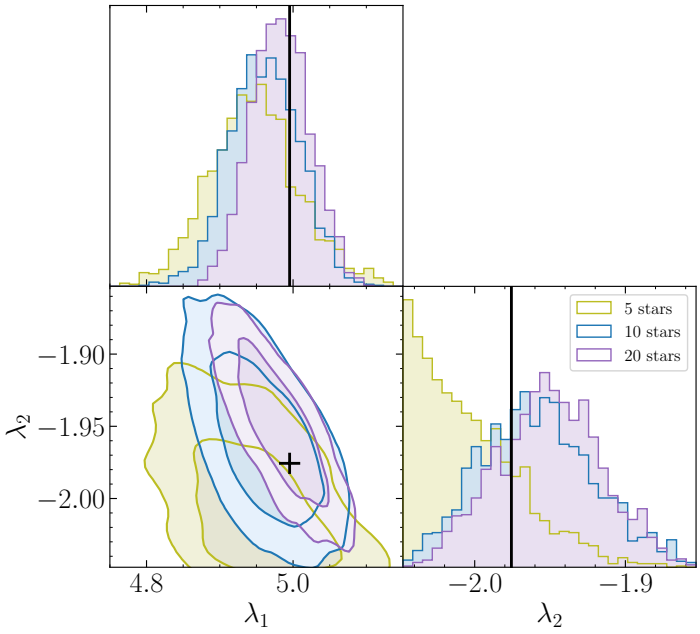

Figure 5. Corner plot depicting the posterior distribution of the parameters λ_1 and λ_2 of one example EoS. The posterior is computed based on the simulated spectra of 5 (olive), 10 (blue), or 20 (purple) stars with the nuisance parameters known with the uncertainties in table 2 of the *loose* scenario. The ground-truth parameter values are depicted as black crosses/lines.

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.

5.3 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 section 3, 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 section 4. From the marginal posterior distributions, similar to the example in figure 2, we determine the maximum-a-posteriori estimates (MAP)[3] for the two EoS parameters $(\lambda_{1,\mathrm{pred}}, \lambda_{2,\mathrm{pred}})$ and compare them to the ground-truth values $(\lambda_{1,\mathrm{truth}}, \lambda_{2,\mathrm{truth}})$. The distributions of the differences between the marginal MAP estimates and the ground-truth values, $(\lambda_{1,\mathrm{pred}} - \lambda_{1,\mathrm{truth}}, \lambda_{2,\mathrm{pred}} - \lambda_{2,\mathrm{truth}})$, on the test data set are depicted in figure 6. As before, the equation of state parameters can be maximally constrained in the true nuisance parameter scenario. In all three scenarios, the distributions are centered

[3]Note that this is the maximum-a-posteriori estimate of the marginal and not of the full posterior.

JCAP09(2024)009

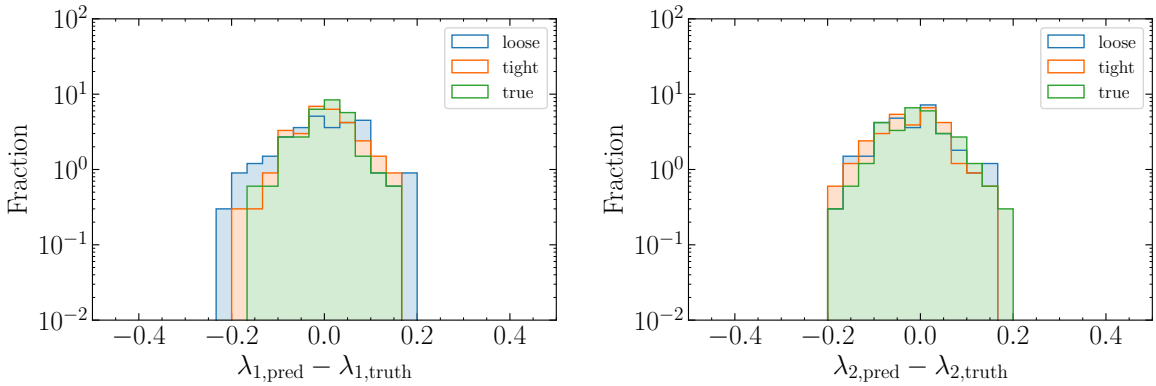

Figure 6. Distribution of the predicted maximum-a-posteriori estimates versus the ground-truth values for the two equation of state parameters λ_1 and λ_2. In the true scenario, the nuisance parameters are fixed to their exact values; in the tight and loose cases, they are drawn from the narrow or wide priors in table 2.

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 λ_1 are much narrower in the tight and loose scenarios (see for example figure 8 in [72]).

To quantitatively compare these distributions to the previous ML analyses, we compute the mean μ and standard deviation σ of the distribution of differences in figure 6. We can combine both standard deviations into

$$\sigma_{\mathrm{tot}} = \sqrt{\sigma_{\lambda_1}^2 + \sigma_{\lambda_2}^2}\,. \tag{5.1}$$

The resulting values are listed in table 3 and illustrated in figure 7. The mean μ of the difference between the posterior prediction of λ_1 and its ground-truth value is much smaller in the NLE approach compared to previous approaches in all three prior distributions considered in table 2. In the true case, the NLE approach performs better than ML-Likelihood$_{\mathrm{EOS}}$ from ref. [72] and NN(Spectra) from ref. [71] (see section 3 for more details about these methods). For realistic scenarios of nuisance parameters (loose and tight), NLE outperforms all other methods as quantified by σ_{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 table 3 it becomes clear that the NLE approach is better than all other approaches to constrain the first EoS parameter λ_1, while for λ_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. Note that the central enthalpies used to solve the stellar structure equations are sampled from log-uniform intervals, such that the masses near the maximum supported mass are weighted more heavily for each EoS. Consequently, for each EoS in the test set, the largest of the ten randomly selected masses is close to the respective maximum supported mass. This is also true for the previous approaches with which we compare our results.

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

JCAP09(2024)009

		$\lambda_{1,\text{pred}} - \lambda_{1,\text{truth}}$		$\lambda_{2,\text{pred}} - \lambda_{2,\text{truth}}$		Combined
$p(\nu)$	Method	μ	σ	μ	σ	σ_{tot}
true	ML-Likelihood$_{\text{EOS}}$	−0.02	0.066	**0.01**	0.070	0.096
	NN(Spectra)	−0.02	0.066	**0.01**	0.075	0.099
	NN(M, R via XSPEC)	−0.03	0.065	**0.01**	**0.055**	**0.085**
	NLE	**0.00**	**0.056**	**−0.01**	0.070	0.090
tight	ML-Likelihood$_{\text{EOS}}$	−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	**0.01**	**0.056**	0.098
	NLE	**0.00**	**0.066**	−0.02	0.071	**0.097**
loose	ML-Likelihood$_{\text{EOS}}$	−0.04	0.089	0.03	0.081	0.120
	NN(Spectra)	−0.03	0.131	**−0.01**	0.078	0.152
	NN(M, R via XSPEC)	−0.03	0.123	**0.01**	**0.058**	0.136
	NLE	**0.00**	**0.085**	**−0.01**	0.074	**0.113**

Table 3. Average accuracy for the prediction of neutron star EoS parameters λ_1 and λ_2. Shown are the means (μ) and standard deviations (σ) of the distributions in figure 6, i.e., of the differences between the predicted maximum-a-posteriori and ground-truth values. Both standard deviations are combined to σ_{tot} according to eq. (5.1). 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$_{\text{EOS}}$, 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 table 2.

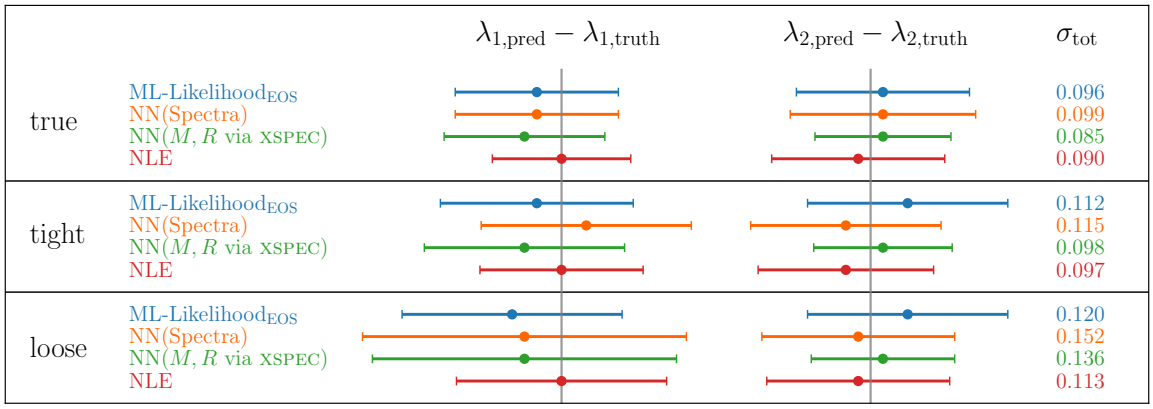

Figure 7. Illustrated mean and standard deviation of the difference between the predicted maximum-a-posteriori values to the ground-truth values for the three different scenarios from table 3.

between the marginal MAP estimates and the ground-truth values on the test data are listed in table 4. 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

JCAP09(2024)009

		$N_{H,\text{pred}} - N_{H,\text{truth}}$		$d_\text{pred} - d_\text{truth}$		$\log(T_\text{eff})_\text{pred} - \log(T_\text{eff})_\text{truth}$	
Method	$p(\nu)$	μ	σ	μ	σ	μ	σ
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

Table 4. Similar to table 3, but with mean and standard deviation of the predicted maximum-a-posteriori estimate versus the ground-truth value for the three nuisance parameters.

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 figure 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.

6 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 [111] can impact a star's spectrum but are not captured by the mass and radius. Uncertainties from nuisance parameters can also be more accurately propagated by using the full high-dimensional data. 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.

(ii) Neural likelihood estimation allows for amortization; after training the neural density estimators *once*, the inclusion of additional observations is straightforward, see section 5.2. 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 [112], e.g., constraints from low-energy nuclear theory at small densities [113, 114], perturbative QCD at high densities [115, 116], 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]. If these likelihoods are computed by traditional methods, they may not be differentiable, which precludes the use of HMC sampling methods, however it may be worthwhile extending simulation-based inference techniques also to these cases.

(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.

JCAP09(2024)009

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. Moreover, to combine the information from different observational instruments for the same neutron star, it becomes essential to compute the likelihood as a function of the neutron star mass. Such situations arise regularly, see e.g. [26, 117–120]. Preliminary results show, however, that the mass distributions can become multimodal, as also seen in [26], which makes sampling the posterior difficult. Future work may extend our method to effectively include the mass as a nuisance parameter. 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 [121], segment-wise linear interpolations of the speed of sound [122], non-parametric representations based on a Gaussian process [37, 123] or a neural network [124]. In that way, we could constrain the unknown phase structure of strongly interacting matter at large densities [125]. Our approach is conducive to the use of a more complex EoS parameterization.

7 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 $\log(T_{\rm eff})$ 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 [126].[4] 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

[4]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 [103].

JCAP09(2024)009

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, 112, 127]. Neural simulation-based inference techniques promise to provide a cost-efficient alternative [92, 128].

This approach can help guide decisions about future observations. Section 5.2 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 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.

Acknowledgments

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 (U.S.A.) 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.

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

JCAP09(2024)009

a scalar meson (σ) signifying attraction between baryons, a vector meson (ω) representing repulsion, and an isovector meson (ρ) 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 ψ_N is expressed as [129]:

$$\mathcal{L}_{\rm N} = \sum_N \bar{\psi}_N \left[\gamma_\mu \left(i\partial^\mu - g_{\omega N}\omega^\mu - \frac{1}{2} g_{\rho N}\, \boldsymbol{\tau}\cdot\boldsymbol{\rho}^\mu \right) - \left(m_N - g_{\sigma N}\,\sigma\right) \right] \psi_N \,, \tag{A.1}$$

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 $n_0 = 0.153\,\mathrm{fm}^{-3}$, $E_0 = -16.3\,\mathrm{MeV}$, $K_0 = 300\,\mathrm{MeV}$, $m^*/m_N = 0.70$, $J = 32.5\,\mathrm{MeV}$, $L_0 = 55.0\,\mathrm{MeV}$, and $U_N = -65.5\,\mathrm{MeV}$. Electrons and muons in neutron star matter are described by the lepton Lagrangian ($\lambda \in \{e^-, \mu^-\}$):

$$\mathcal{L}_{\rm L} = \sum_\lambda \overline{\psi}_\lambda (i\gamma_\mu \partial^\mu - m_\lambda)\psi_\lambda \,, \tag{A.2}$$

while the meson Lagrangian is given by [129, 130]:

$$\mathcal{L}_{\rm M} = \frac{1}{2}(\partial_\mu\sigma\partial^\mu\sigma - m_\sigma^2\sigma^2) - \frac{1}{4}\omega_{\mu\nu}\omega^{\mu\nu} + \frac{1}{2}m_\omega^2\omega_\mu\omega^\mu + \frac{1}{2}m_\rho^2\boldsymbol{\rho}_\mu\cdot\boldsymbol{\rho}^\mu - \frac{1}{4}\boldsymbol{\rho}_{\mu\nu}\cdot\boldsymbol{\rho}^{\mu\nu} \,, \tag{A.3}$$

where $\omega_{\mu\nu} = \partial_\mu\omega_\nu - \partial_\nu\omega_\mu$ and $\boldsymbol{\rho}_{\mu\nu} = \partial_\mu\boldsymbol{\rho}_\nu - \partial_\nu\boldsymbol{\rho}_\mu$. 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 [131]:

$$\mathcal{L}_{{\rm NL}\sigma} = -\frac{1}{3}b_\sigma m_N \left[g_{\sigma N}(n)\sigma\right]^3 - \frac{1}{4}c_\sigma \left[g_{\sigma N}(n)\sigma\right]^4 , \tag{A.4}$$

where b_σ and c_σ 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:

$$n\,q_p + \sum_\lambda n_\lambda q_\lambda = 0 \,, \tag{A.5}$$

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

$$n - \sum_{B=n,p} n_B = 0 \,. \tag{A.6}$$

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

JCAP09(2024)009

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:

$$\mu_p = \mu_n - q_p \mu_e \,, \tag{A.7}$$

where μ_n and μ_p are the chemical potentials of neutrons and protons respectively, and μ_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 [132] for the outer crust and the Baym-Bethe-Pethick (BBP) model [133] 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 λ_1 and λ_2. The original coefficients describing the GM1L EoS are then used to generate many EoS scenarios using the expression:

$$\lambda_{\text{generated}} = \lambda_{\text{fit}} \cdot (1 + 2 \cdot c\,(ran2 - 0.5)) \tag{A.8}$$

where $\lambda_{\text{generated}}$ represents the newly constructed spectral parameter, λ_{fit} is the best fit spectral parameter of GM1L, and c is a scaling parameter set to 0.05. *ran2* are uniformly distributed random numbers in the range 0 to 1 generated by the *ran2* function given in [134]. This process was repeated to create around 10^3 different EoS models, each represented by a unique set of spectral coefficients λ_1 and λ_2. Following this process, the two parameters are uniformly distributed in the intervals $\lambda_1 \in [4.75, 5.25]$ and $\lambda_2 \in [-2.05, -1.85]$.

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_Φ with trainable parameters Φ.

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

$$u \sim \pi(u)\,, \tag{B.1}$$

$$x = f_\Phi(u)\,. \tag{B.2}$$

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

$$p(x) = \pi(f_\Phi^{-1}(x)) \left| \det \left(\frac{\partial f_\Phi^{-1}}{\partial x} \right) \right| . \tag{B.3}$$

JCAP09(2024)009

Consequently, for a fast numerical evaluation of the probability density, the transformations f_Φ 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) [102], where each dimension of the data is sequentially transformed conditioned on the previous dimensions, making the Jacobian of f_Φ^{-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 my minimizing the Kullback-Leibler divergence ($D_{\rm KL}$), 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, $q_\Phi(s|\lambda,\nu)$. $D_{\rm KL}$ becomes zero if the distributions are identical. Hence we fit the trainable parameters Φ with the following optimization procedure:

$$\begin{aligned}
\arg\min_\Phi D_{\rm KL}\left(p(s|\lambda,\nu)||q_\Phi(s|\lambda,\nu)\right) &= \arg\min_\Phi \int \mathrm{d}s\, p(s|\lambda,\nu)\left[\log p(s|\lambda,\nu) - \log q_\Phi(s|\lambda,\nu)\right] \\
&\approx \arg\min_\Phi \sum_{s_i\sim p(s|\lambda_i,\nu_i)} \log p(s_i|\lambda_i,\nu_i) - \log q_\Phi(s_i|\lambda_i,\nu_i) \\
&= \arg\min_\Phi \sum_{s_i\sim p(s|\lambda_i,\nu_i)} -\log q_\Phi(s_i|\lambda_i,\nu_i) \\
&= \arg\max_\Phi \sum_{s_i\sim p(s|\lambda_i,\nu_i)} \log q_\Phi(s_i|\lambda_i,\nu_i)\,, \qquad \text{(B.4)}
\end{aligned}$$

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 Φ 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, (λ_i,ν_i).

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.

C.1 Metropolis-Hastings

Given the target distribution π 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 σ, 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

JCAP09(2024)009

Markov chain generates samples from the target distribution, $x, y \sim \pi$. This is equivalent to ensuring that the target distribution π is invariant under this Markov chain, for example by maintaining detailed balance conditions when accepting the proposed stage

$$\pi(x)t(x,y)\alpha(x,y) = \pi(y)t(y,x)\alpha(y,x)\,. \tag{C.1}$$

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 ($x \to y$) is the same as making the reverse transition ($y \to x$).

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

$$\alpha(x,y) = \min\left(\frac{\pi(y)\,t(y,x)}{\pi(x)\,t(x,y)}, 1\right), \tag{C.2}$$

where the denominator is never zero given the above assumptions on π and t. For each $x, y \in \mathcal{R}^n$, either $\alpha(x,y)$ or $\alpha(y,x)$ is 1. There are other formulae for α obeying eq. (C.1), but with lower acceptance rates, meaning they lead to an undesirable higher asymptotic variance for estimated expectations.

C.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 [105, 106]. Following standard notation for Hamiltonian dynamics, $q \in \mathcal{R}^n$ denotes the parameter of interest that is to be sampled. The target density, π, 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 $H : \mathcal{R}^{2n} \to \mathcal{R}$ gives the total energy for the state $x := (q,p)$,

$$H(x) = H(q,p) = U(q) + \frac{1}{2}p^T M^{-1} p\,. \tag{C.3}$$

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

$$\tilde{\pi}(x) := Z^{-1}\, e^{-H(x)} = Z^{-1}\, e^{-U(q)}\, e^{-\frac{1}{2}p^T M^{-1} p} = Z^{-1}\, \pi(q)\, e^{-\frac{1}{2}p^T M^{-1} p}\,, \tag{C.4}$$

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

JCAP09(2024)009

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 π.

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. (C.3) over a certain time T, starting from the initial point $x = (q^{(i)}, p^{(i)})$ and is followed by negation of the final momentum.[5] 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:

$$\begin{aligned} \bar{p} &\leftarrow p - \frac{\epsilon}{2}\nabla U(q)\,, \\ q' &\leftarrow q + \epsilon M^{-1}\bar{p}\,, \\ p' &\leftarrow \bar{p} - \frac{\epsilon}{2}\nabla U(q')\,, \end{aligned} \tag{C.5}$$

where ϵ 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 $P(q, p) := (q, -p)$, $F = (L_\epsilon)^n P$ (operators act left to right) is volume-preserving because each of the three substeps is a shear transformation,[6] and F is an involution[7] 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

$$\alpha(x, y) = \min\left(\frac{\tilde{\pi}(y)}{\tilde{\pi}(x)}, 1\right), \tag{C.6}$$

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.

[5] 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.

[6] 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.

[7] Involution means $F^{-1} = F$, i.e., it is *time reversible*.

JCAP09(2024)009

C.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 [135]. We use an importance sampling step to determine initial values for each chain [90]. To facilitate an easy comparison of our NLE + HMC method with future analyses, we provide working examples of the algorithms implemented for this work in a public GitHub repository.[8]

References

[1] G. Baym, T. Hatsuda, T. Kojo, P.D. Powell, Y. Song and T. Takatsuka, *From hadrons to quarks in neutron stars: a review*, *Rept. Prog. Phys.* **81** (2018) 056902 [arXiv:1707.04966] [INSPIRE].

[2] D. Blaschke and N. Chamel, *Phases of dense matter in compact stars*, in *The Physics and Astrophysics of Neutron Stars*, L. Rezzolla, P. Pizzochero, D.I. Jones, N. Rea and I. Vidaña eds., *Astrophysics and Space Science Library* **457**, Springer, Cham, Switzerland (2018), pp. 337–400 [DOI:10.1007/978-3-319-97616-7_7] [arXiv:1803.01836] [INSPIRE].

[3] L. Tolos and L. Fabbietti, *Strangeness in Nuclei and Neutron Stars*, *Prog. Part. Nucl. Phys.* **112** (2020) 103770 [arXiv:2002.09223] [INSPIRE].

[4] J.-J. Li, A. Sedrakian and F. Weber, *Competition between delta isobars and hyperons and properties of compact stars*, *Phys. Lett. B* **783** (2018) 234 [arXiv:1803.03661] [INSPIRE].

[5] W.M. Spinella and F. Weber, *Hyperonic Neutron Star Matter in Light of GW170817*, *Astron. Nachr.* **340** (2019) 145 [arXiv:1812.03600] [INSPIRE].

[6] G. Malfatti, M.G. Orsaria, I.F. Ranea-Sandoval, G.A. Contrera and F. Weber, *Delta baryons and diquark formation in the cores of neutron stars*, *Phys. Rev. D* **102** (2020) 063008 [arXiv:2008.06459] [INSPIRE].

[7] A. Sedrakian, J.-J. Li and F. Weber, *Heavy baryons in compact stars*, *Prog. Part. Nucl. Phys.* **131** (2023) 104041 [arXiv:2212.01086] [INSPIRE].

[8] C. Alcock, E. Farhi and A. Olinto, *Strange stars*, *Astrophys. J.* **310** (1986) 261 [INSPIRE].

[9] J. Madsen, *Physics and Astrophysics of Strange Quark Matter*, in *Hadrons in Dense Matter and Hadrosynthesis*, *Lecture Notes in Physics* **516**, J. Cleymans, H.B. Geyer and F.G. Scholtz eds., Springer (1999), pp. 162–203 [astro-ph/9809032] [DOI:10.1007/BFb0107314].

[10] F. Weber, *Strange quark matter and compact stars*, *Prog. Part. Nucl. Phys.* **54** (2005) 193 [astro-ph/0407155] [INSPIRE].

[11] M. Orsaria, H. Rodrigues, F. Weber and G.A. Contrera, *Quark deconfinement in high-mass neutron stars*, *Phys. Rev. C* **89** (2014) 015806 [arXiv:1308.1657] [INSPIRE].

[12] M.G. Alford, *Color superconducting quark matter*, *Annu. Rev. Nucl. Part. Sci.* **51** (2001) 131 [hep-ph/0102047] [INSPIRE].

[8] https://github.com/lenjonah/neutron_star_inference.

JCAP09(2024)009

[13] M.G. Alford, A. Schmitt, K. Rajagopal and T. Schäfer, *Color superconductivity in dense quark matter*, *Rev. Mod. Phys.* **80** (2008) 1455 [arXiv:0709.4635] [INSPIRE].

[14] J.L. Zdunik and P. Haensel, *Maximum mass of neutron stars and strange neutron-star cores*, *Astron. Astrophys.* **551** (2013) A61 [arXiv:1211.1231] [INSPIRE].

[15] M. Buballa, *NJL model analysis of quark matter at large density*, *Phys. Rep.* **407** (2005) 205 [hep-ph/0402234] [INSPIRE].

[16] L. McLerran and S. Reddy, *Quarkyonic Matter and Neutron Stars*, *Phys. Rev. Lett.* **122** (2019) 122701 [arXiv:1811.12503] [INSPIRE].

[17] G. Baym, *Pion condensation in nuclear and neutron star matter*, *Phys. Rev. Lett.* **30** (1973) 1340 [INSPIRE].

[18] D.B. Kaplan and A.E. Nelson, *Strange Goings on in Dense Nucleonic Matter*, *Phys. Lett. B* **175** (1986) 57 [INSPIRE].

[19] N.K. Glendenning and J. Schaffner-Bielich, *First order kaon condensate*, *Phys. Rev. C* **60** (1999) 025803 [astro-ph/9810290] [INSPIRE].

[20] P.J. Ellis, R. Knorren and M. Prakash, *Kaon condensation in neutron star matter with hyperons*, *Phys. Lett. B* **349** (1995) 11 [nucl-th/9502033] [INSPIRE].

[21] A. Ramos, J. Schaffner-Bielich and J. Wambach, *Kaon Condensation in Neutron Stars*, in *Physics of Neutron Star Interiors*, D. Blaschke, A. Sedrakian and N.K. Glendenning eds., Springer (2001), pp. 175–202 [DOI:10.1007/3-540-44578-1_6].

[22] T. Kojo, *QCD equations of state and speed of sound in neutron stars*, *AAPPS Bull.* **31** (2021) 11 [arXiv:2011.10940] [INSPIRE].

[23] R.C. Tolman, *Static solutions of Einstein's field equations for spheres of fluid*, *Phys. Rev.* **55** (1939) 364 [INSPIRE].

[24] J.R. Oppenheimer and G.M. Volkoff, *On massive neutron cores*, *Phys. Rev.* **55** (1939) 374 [INSPIRE].

[25] J.M. Lattimer and A.W. Steiner, *Neutron Star Masses and Radii from Quiescent Low-Mass X-ray Binaries*, *Astrophys. J.* **784** (2014) 123 [arXiv:1305.3242] [INSPIRE].

[26] A.W. Steiner et al., *Constraining the Mass and Radius of Neutron Stars in Globular Clusters*, *Mon. Not. Roy. Astron. Soc.* **476** (2018) 421 [arXiv:1709.05013] [INSPIRE].

[27] S. Ascenzi, V. Graber and N. Rea, *Neutron-star measurements in the multi-messenger Era*, *Astropart. Phys.* **158** (2024) 102935 [arXiv:2401.14930] [INSPIRE].

[28] J. Antoniadis et al., *A Massive Pulsar in a Compact Relativistic Binary*, *Science* **340** (2013) 6131 [arXiv:1304.6875] [INSPIRE].

[29] Z. Arzoumanian et al., *The NANOGrav* 11*-year Data Set: High-precision timing of* 45 *Millisecond Pulsars*, *Astrophys. J. Suppl. Ser.* **235** (2018) 37 [arXiv:1801.01837] [INSPIRE].

[30] E. Fonseca et al., *Refined Mass and Geometric Measurements of the High-mass PSR J0740+6620*, *Astrophys. J. Lett.* **915** (2021) L12 [arXiv:2104.00880] [INSPIRE].

[31] T.E. Riley et al., *A NICER View of PSR J0030+0451: Millisecond Pulsar Parameter Estimation*, *Astrophys. J. Lett.* **887** (2019) L21 [arXiv:1912.05702] [INSPIRE].

[32] M.C. Miller et al., *The Radius of PSR J0740+6620 from NICER and XMM-Newton Data*, *Astrophys. J. Lett.* **918** (2021) L28 [arXiv:2105.06979] [INSPIRE].

JCAP09(2024)009

[33] M.C. Miller et al., *PSR J0030+0451 Mass and Radius from NICER Data and Implications for the Properties of Neutron Star Matter*, *Astrophys. J. Lett.* **887** (2019) L24 [arXiv:1912.05705] [INSPIRE].

[34] LIGO Scientific and Virgo collaborations, *Properties of the binary neutron star merger GW170817*, *Phys. Rev. X* **9** (2019) 011001 [arXiv:1805.11579] [INSPIRE].

[35] LIGO Scientific and Virgo collaborations, *GW190425: Observation of a Compact Binary Coalescence with Total Mass* $\sim 3.4\,M_{\odot}$, *Astrophys. J. Lett.* **892** (2020) L3 [arXiv:2001.01761] [INSPIRE].

[36] G. Raaijmakers et al., *A NICER view of PSR J0030+0451: Implications for the dense matter equation of state*, *Astrophys. J. Lett.* **887** (2019) L22 [arXiv:1912.05703] [INSPIRE].

[37] I. Legred, K. Chatziioannou, R. Essick, S. Han and P. Landry, *Impact of the PSR J0740+6620 radius constraint on the properties of high-density matter*, *Phys. Rev. D* **104** (2021) 063003 [arXiv:2106.05313] [INSPIRE].

[38] L. Brandes, W. Weise and N. Kaiser, *Inference of the sound speed and related properties of neutron stars*, *Phys. Rev. D* **107** (2023) 014011 [arXiv:2208.03026] [INSPIRE].

[39] L. Brandes, W. Weise and N. Kaiser, *Evidence against a strong first-order phase transition in neutron star cores: Impact of new data*, *Phys. Rev. D* **108** (2023) 094014 [arXiv:2306.06218] [INSPIRE].

[40] S. Huth et al., *Constraining Neutron-Star Matter with Microscopic and Macroscopic Collisions*, *Nature* **606** (2022) 276 [arXiv:2107.06229] [INSPIRE].

[41] S. Altiparmak, C. Ecker and L. Rezzolla, *On the Sound Speed in Neutron Stars*, *Astrophys. J. Lett.* **939** (2022) L34 [arXiv:2203.14974] [INSPIRE].

[42] Y. Lim and J.W. Holt, *Neutron Star Radii, Deformabilities, and Moments of Inertia from Experimental and Ab Initio Theory Constraints of the* ^{208}Pb *Neutron Skin Thickness*, *Galaxies* **10** (2022) 99 [arXiv:2204.09000] [INSPIRE].

[43] M. Marczenko, L. McLerran, K. Redlich and C. Sasaki, *Reaching percolation and conformal limits in neutron stars*, *Phys. Rev. C* **107** (2023) 025802 [arXiv:2207.13059] [INSPIRE].

[44] E. Annala et al., *Strongly interacting matter exhibits deconfined behavior in massive neutron stars*, *Nat. Commun.* **14** (2023) 8451 [arXiv:2303.11356] [INSPIRE].

[45] M.-Z. Han, Y.-J. Huang, S.-P. Tang and Y.-Z. Fan, *Plausible presence of new state in neutron stars with masses above 0.98MTOV*, *Sci. Bull.* **68** (2023) 913 [arXiv:2207.13613] [INSPIRE].

[46] J.-L. Jiang, C. Ecker and L. Rezzolla, *Bayesian Analysis of Neutron-star Properties with Parameterized Equations of State: The Role of the Likelihood Functions*, *Astrophys. J.* **949** (2023) 11 [arXiv:2211.00018] [INSPIRE].

[47] R. Essick, I. Legred, K. Chatziioannou, S. Han and P. Landry, *Phase transition phenomenology with nonparametric representations of the neutron star equation of state*, *Phys. Rev. D* **108** (2023) 043013 [arXiv:2305.07411] [INSPIRE].

[48] D. Mroczek, M.C. Miller, J. Noronha-Hostler and N. Yunes, *Nontrivial features in the speed of sound inside neutron stars*, arXiv:2309.02345 [INSPIRE].

[49] P.T.H. Pang et al., *Probing quarkyonic matter in neutron stars with the Bayesian nuclear-physics multimessenger astrophysics framework*, *Phys. Rev. C* **109** (2024) 025807 [arXiv:2308.15067] [INSPIRE].

JCAP09(2024)009

[50] H. Koehn et al., *An overview of existing and new nuclear and astrophysical constraints on the equation of state of neutron-rich dense matter*, arXiv:2402.04172 [INSPIRE].

[51] Y. Fujimoto, K. Fukushima and K. Murase, *Methodology study of machine learning for the neutron star equation of state*, *Phys. Rev. D* **98** (2018) 023019 [arXiv:1711.06748] [INSPIRE].

[52] Y. Fujimoto, K. Fukushima and K. Murase, *Mapping neutron star data to the equation of state using the deep neural network*, *Phys. Rev. D* **101** (2020) 054016 [arXiv:1903.03400] [INSPIRE].

[53] Y. Fujimoto, K. Fukushima and K. Murase, *Extensive Studies of the Neutron Star Equation of State from the Deep Learning Inference with the Observational Data Augmentation*, *JHEP* **03** (2021) 273 [arXiv:2101.08156] [INSPIRE].

[54] M. Ferreira and C. Providência, *Unveiling the nuclear matter EoS from neutron star properties: a supervised machine learning approach*, *JCAP* **07** (2021) 011 [arXiv:1910.05554] [INSPIRE].

[55] F. Morawski and M. Bejger, *Neural network reconstruction of the dense matter equation of state derived from the parameters of neutron stars*, *Astron. Astrophys.* **642** (2020) A78 [arXiv:2006.07194] [INSPIRE].

[56] S. Traversi and P. Char, *Structure of Quark Star: A Comparative Analysis of Bayesian Inference and Neural Network Based Modeling*, *Astrophys. J.* **905** (2020) 9 [arXiv:2007.10239] [INSPIRE].

[57] P.G. Krastev, *Translating Neutron Star Observations to Nuclear Symmetry Energy via Deep Neural Networks*, *Galaxies* **10** (2022) 16 [arXiv:2112.04089] [INSPIRE].

[58] F. Morawski and M. Bejger, *Detecting dense-matter phase transition signatures in neutron star mass-radius measurements as data anomalies using normalizing flows*, *Phys. Rev. C* **106** (2022) 065802 [arXiv:2212.05480] [INSPIRE].

[59] M. Ferreira, V. Carvalho and C. Providência, *Extracting nuclear matter properties from the neutron star matter equation of state using deep neural networks*, *Phys. Rev. D* **106** (2022) 103023 [arXiv:2209.09085] [INSPIRE].

[60] A. Thete, K. Banerjee and T. Malik, *Realizing the potential of deep neural network for analyzing neutron star observables and dense matter equation of state*, *Phys. Rev. D* **108** (2023) 063028 [arXiv:2208.13163] [INSPIRE].

[61] S. Soma, L. Wang, S. Shi, H. Stöcker and K. Zhou, *Neural network reconstruction of the dense matter equation of state from neutron star observables*, *JCAP* **08** (2022) 071 [arXiv:2201.01756] [INSPIRE].

[62] S. Soma, L. Wang, S. Shi, H. Stöcker and K. Zhou, *Reconstructing the neutron star equation of state from observational data via automatic differentiation*, *Phys. Rev. D* **107** (2023) 083028 [arXiv:2209.08883] [INSPIRE].

[63] P.G. Krastev, *A Deep Learning Approach to Extracting Nuclear Matter Properties from Neutron Star Observations*, *Symmetry* **15** (2023) 1123 [arXiv:2303.17146] [INSPIRE].

[64] L.-J. Guo, J.-Y. Xiong, Y. Ma and Y.-L. Ma, *Insights into Neutron Star Equation of State by Machine Learning*, *Astrophys. J.* **965** (2024) 47 [arXiv:2309.11227] [INSPIRE].

[65] S. Chatterjee, H. Sudhakaran and R. Mallick, *Analyzing the speed of sound in neutron star with machine learning*, arXiv:2302.13648 [INSPIRE].

[66] W. Zhou, J. Hu, Y. Zhang and H. Shen, *Nonparametric Model for the Equations of State of a Neutron Star from Deep Neural Network*, *Astrophys. J.* **950** (2023) 186 [arXiv:2305.03323] [INSPIRE].

JCAP09(2024)009

[67] V. Carvalho, M. Ferreira, T. Malik and C. Providência, *Decoding neutron star observations: Revealing composition through Bayesian neural networks*, *Phys. Rev. D* **108** (2023) 043031 [arXiv:2306.06929] [INSPIRE].

[68] V. Carvalho, M. Ferreira and C. Providência, *From neutron star observations to nuclear matter properties: A machine learning approach*, *Phys. Rev. D* **109** (2024) 123038 [arXiv:2401.05770] [INSPIRE].

[69] Y. Fujimoto, K. Fukushima, S. Kamata and K. Murase, *Uncertainty quantification in the machine-learning inference from neutron star probability distribution to the equation of state*, arXiv:2401.12688 [INSPIRE].

[70] K. Zhou, L. Wang, L.-G. Pang and S. Shi, *Exploring QCD matter in extreme conditions with Machine Learning*, *Prog. Part. Nucl. Phys.* **135** (2024) 104084 [arXiv:2303.15136] [INSPIRE].

[71] D. Farrell et al., *Deducing neutron star equation of state parameters directly from telescope spectra with uncertainty-aware machine learning*, *JCAP* **02** (2023) 016 [arXiv:2209.02817] [INSPIRE].

[72] D. Farrell et al., *Deducing neutron star equation of state from telescope spectra with machine-learning-derived likelihoods*, *JCAP* **12** (2023) 022 [arXiv:2305.07442] [INSPIRE].

[73] T.E. Riley, G. Raaijmakers and A.L. Watts, *On parametrized cold dense matter equation-of-state inference*, *Mon. Not. Roy. Astron. Soc.* **478** (2018) 1093 [arXiv:1804.09085] [INSPIRE].

[74] M.C. Miller, C. Chirenti and F.K. Lamb, *Constraining the equation of state of high-density cold matter using nuclear and astronomical measurements*, *Astrophys. J.* **888** (2019) 12 [arXiv:1904.08907] [INSPIRE].

[75] G. Papamakarios, D.C. Sterratt and I. Murray, *Sequential Neural Likelihood: Fast Likelihood-free Inference with Autoregressive Flows*, arXiv:1805.07226.

[76] L. Dinh, J. Sohl-Dickstein and S. Bengio, *Density estimation using Real NVP*, arXiv:1605.08803 [INSPIRE].

[77] G. Papamakarios, E. Nalisnick, D.J. Rezende, S. Mohamed and B. Lakshminarayanan, *Normalizing Flows for Probabilistic Modeling and Inference*, *J. Mach. Learn. Res.* **22** (2021) 1 [arXiv:1912.02762] [INSPIRE].

[78] S. Typel, G. Röpke, T. Klähn, D. Blaschke and H.H. Wolter, *Composition and thermodynamics of nuclear matter with light clusters*, *Phys. Rev. C* **81** (2010) 015803 [arXiv:0908.2344] [INSPIRE].

[79] L. Lindblom, *Spectral Representations of Neutron-Star Equations of State*, *Phys. Rev. D* **82** (2010) 103011 [arXiv:1009.0738] [INSPIRE].

[80] L. Lindblom, *Causal Representations of Neutron-Star Equations of State*, *Phys. Rev. D* **97** (2018) 123019 [arXiv:1804.04072] [INSPIRE].

[81] L. Lindblom, *Determining the nuclear equation of state from neutron-star masses and radii*, *Astrophys. J.* **398** (1992) 569.

[82] M.C. Miller and F.K. Lamb, *Observational Constraints on Neutron Star Masses and Radii*, *Eur. Phys. J. A* **52** (2016) 63 [arXiv:1604.03894] [INSPIRE].

[83] M.C. Miller, *Astrophysical Constraints on Dense Matter in Neutron Stars*, *Astrophysics and Space Science Library* **461** (2020) 1 [arXiv:1312.0029] [INSPIRE].

JCAP09(2024)009

[84] G.B. Rybicki, C.O. Heinke, R. Narayan and J.E. Grindlay, *A Hydrogen atmosphere spectral model applied to the neutron star X7 in the globular cluster 47 Tucanae*, *Astrophys. J.* **644** (2006) 1090 [astro-ph/0506563] [INSPIRE].

[85] K. Arnaud, *XSPEC: The First Ten Years*, in *Astronomical Data Analysis Software and Systems V* **101**, *Astronomical Society of the Pacific Conference Series*, Astronomical Society of the Pacific (1996), pp. 17–20.

[86] S. Bogdanov, C.O. Heinke, F. Özel and T. Güver, *Neutron Star Mass-Radius Constraints of the Quiescent Low-mass X-ray Binaries X7 and X5 in the Globular Cluster 47 Tuc*, *Astrophys. J.* **831** (2016) 184 [arXiv:1603.01630] [INSPIRE].

[87] B.W. Silverman, *Density estimation for statistics and data analysis*, Chapman and Hall, London, U.K. (1986).

[88] G. Raaijmakers et al., *Constraints on the Dense Matter Equation of State and Neutron Star Properties from NICER's Mass-Radius Estimate of PSR J0740+6620 and Multimessenger Observations*, *Astrophys. J. Lett.* **918** (2021) L29 [arXiv:2105.06981] [INSPIRE].

[89] K. Cranmer, J. Brehmer and G. Louppe, *The frontier of simulation-based inference*, *Proc. Natl. Acad. Sci. U.S.A.* **117** (2020) 30055 [arXiv:1911.01429] [INSPIRE].

[90] J.-M. Lueckmann, J. Boelts, D.S. Greenberg, P.J. Gonçalves and J.H. Macke, *Benchmarking Simulation-Based Inference*, arXiv:2101.04653.

[91] A. Delaunoy et al., *Lightning-Fast Gravitational Wave Parameter Inference through Neural Amortization*, arXiv:2010.12931 [INSPIRE].

[92] M. Dax, S.R. Green, J. Gair, J.H. Macke, A. Buonanno and B. Schölkopf, *Real-Time Gravitational Wave Science with Neural Posterior Estimation*, *Phys. Rev. Lett.* **127** (2021) 241103 [arXiv:2106.12594] [INSPIRE].

[93] S.R. Green and J. Gair, *Complete parameter inference for GW150914 using deep learning*, *Mach. Learn. Sci. Technol.* **2** (2021) 03LT01 [arXiv:2008.03312] [INSPIRE].

[94] K. Cranmer, J. Pavez and G. Louppe, *Approximating Likelihood Ratios with Calibrated Discriminative Classifiers*, arXiv:1506.02169 [INSPIRE].

[95] J. Brehmer, F. Kling, I. Espejo and K. Cranmer, *MadMiner: Machine learning-based inference for particle physics*, *Comput. Softw. Big Sci.* **4** (2020) 3 [arXiv:1907.10621] [INSPIRE].

[96] A. Ghosh, *Measuring quantum interference in the off-shell Higgs to flour leptons process with Machine Learning*, in *Journées de Rencontre des Jeunes Chercheurs 2019. Book of Proceedings*, Société Française de Physique (2020), pp. 171–176 [hal-02971995v3].

[97] J. Alsing, T. Charnock, S. Feeney and B. Wandelt, *Fast likelihood-free cosmology with neural density estimators and active learning*, *Mon. Not. Roy. Astron. Soc.* **488** (2019) 4440 [arXiv:1903.00007] [INSPIRE].

[98] N. Jeffrey, J. Alsing and F. Lanusse, *Likelihood-free inference with neural compression of DES SV weak lensing map statistics*, *Mon. Not. Roy. Astron. Soc.* **501** (2021) 954 [arXiv:2009.08459] [INSPIRE].

[99] C.H. Hahn et al., *A forward modeling approach to analyzing galaxy clustering with SimBIG*, *Proc. Natl. Acad. Sci. U.S.A.* **120** (2023) e2218810120 [arXiv:2211.00723] [INSPIRE].

[100] C. Modi and O.H.E. Philcox, *Hybrid SBI or How I Learned to Stop Worrying and Learn the Likelihood*, arXiv:2309.10270 [INSPIRE].

JCAP09(2024)009

[101] C. Modi, S. Pandey, M. Ho, C. Hahn, B.R.-S. Blancard and B. Wandelt, *Sensitivity Analysis of Simulation-Based Inference for Galaxy Clustering*, arXiv:2309.15071 [INSPIRE].

[102] G. Papamakarios, T. Pavlakou and I. Murray, *Masked Autoregressive Flow for Density Estimation*, in *Advances in Neural Information Processing Systems* **30**, Curran Associates, Inc. (2017) [arXiv:1705.07057] [INSPIRE].

[103] A. Tejero-Cantero et al., *sbi: A toolkit for simulation-based inference*, *J. Open Source Softw.* **5** (2020) 2505.

[104] J. Hermans, A. Delaunoy, F. Rozet, A. Wehenkel and G. Louppe, *A Trust Crisis In Simulation-Based Inference? Your Posterior Approximations Can Be Unfaithful*, arXiv:2110.06581 [INSPIRE].

[105] R.M. Neal, *MCMC using Hamiltonian dynamics*, in *Handbook of Markov Chain Monte Carlo*, CRC Press (2011) [arXiv:1206.1901] [DOI:10.1201/b10905] [INSPIRE].

[106] M. Betancourt, *A Conceptual Introduction to Hamiltonian Monte Carlo*, arXiv:1701.02434 [INSPIRE].

[107] G. Papamakarios and I. Murray, *Fast ϵ-free Inference of Simulation Models with Bayesian Conditional Density Estimation*, arXiv:1605.06376 [INSPIRE].

[108] D.S. Greenberg, M. Nonnenmacher and J.H. Macke, *Automatic Posterior Transformation for Likelihood-Free Inference*, arXiv:1905.07488.

[109] C. Durkan, I. Murray and G. Papamakarios, *On Contrastive Learning for Likelihood-free Inference*, arXiv:2002.03712.

[110] J. Hermans, V. Begy and G. Louppe, *Likelihood-free MCMC with Amortized Approximate Ratio Estimators*, arXiv:1903.04057 [INSPIRE].

[111] K.G. Elshamouty, C.O. Heinke, S.M. Morsink, S. Bogdanov and A.L. Stevens, *The impact of surface temperature inhomogeneities on quiescent neutron star radius measurements*, *Astrophys. J.* **826** (2016) 162 [arXiv:1605.09400] [INSPIRE].

[112] J. Lange, R. O'Shaughnessy and M. Rizzo, *Rapid and accurate parameter inference for coalescing, precessing compact binaries*, arXiv:1805.10457 [INSPIRE].

[113] C. Drischler, S. Han and S. Reddy, *Large and massive neutron stars: Implications for the sound speed within QCD of dense matter*, *Phys. Rev. C* **105** (2022) 035808 [arXiv:2110.14896] [INSPIRE].

[114] J. Keller, K. Hebeler and A. Schwenk, *Nuclear Equation of State for Arbitrary Proton Fraction and Temperature Based on Chiral Effective Field Theory and a Gaussian Process Emulator*, *Phys. Rev. Lett.* **130** (2023) 072701 [arXiv:2204.14016] [INSPIRE].

[115] T. Gorda, A. Kurkela, R. Paatelainen, S. Säppi and A. Vuorinen, *Soft Interactions in Cold Quark Matter*, *Phys. Rev. Lett.* **127** (2021) 162003 [arXiv:2103.05658] [INSPIRE].

[116] O. Komoltsev and A. Kurkela, *How Perturbative QCD Constrains the Equation of State at Neutron-Star Densities*, *Phys. Rev. Lett.* **128** (2022) 202701 [arXiv:2111.05350] [INSPIRE].

[117] D. Klochkov, V. Suleimanov, G. Pühlhofer, D.G. Yakovlev, A. Santangelo and K. Werner, *The neutron star in HESSJ1731-347: Central compact objects as laboratories to study the equation of state of superdense matter*, *Astron. Astrophys.* **573** (2015) A53 [arXiv:1410.1055] [INSPIRE].

[118] V. Doroshenko, V. Suleimanov, G. Pühlhofer and A. Santangelo, *A strangely light neutron star within a supernova remnant*, *Nat. Astron.* **6** (2022) 1444.

JCAP09(2024)009

[119] T.E. Riley et al., *A NICER View of the Massive Pulsar PSR J0740+6620 Informed by Radio Timing and XMM-Newton Spectroscopy*, *Astrophys. J. Lett.* **918** (2021) L27 [arXiv:2105.06980] [INSPIRE].

[120] B. Posselt and G.G. Pavlov, *The Cooling of the Central Compact Object in Cas A from 2006 to 2020*, *Astrophys. J.* **932** (2022) 83 [arXiv:2205.06552] [INSPIRE].

[121] S.K. Greif, G. Raaijmakers, K. Hebeler, A. Schwenk and A.L. Watts, *Equation of state sensitivities when inferring neutron star and dense matter properties*, *Mon. Not. Roy. Astron. Soc.* **485** (2019) 5363 [arXiv:1812.08188] [INSPIRE].

[122] E. Annala, T. Gorda, A. Kurkela, J. Nättilä and A. Vuorinen, *Evidence for quark-matter cores in massive neutron stars*, *Nat. Phys.* **16** (2020) 907 [arXiv:1903.09121] [INSPIRE].

[123] P. Landry, R. Essick and K. Chatziioannou, *Nonparametric constraints on neutron star matter with existing and upcoming gravitational wave and pulsar observations*, *Phys. Rev. D* **101** (2020) 123007 [arXiv:2003.04880] [INSPIRE].

[124] M.-Z. Han, J.-L. Jiang, S.-P. Tang and Y.-Z. Fan, *Bayesian Nonparametric Inference of the Neutron Star Equation of State via a Neural Network*, *Astrophys. J.* **919** (2021) 11 [arXiv:2103.05408] [INSPIRE].

[125] L. Brandes and W. Weise, *Constraints on Phase Transitions in Neutron Star Matter*, *Symmetry* **16** (2024) 111 [arXiv:2312.11937] [INSPIRE].

[126] G. Gonçalves et al., *Machine-learning Love: classifying the equation of state of neutron stars with transformers*, *JCAP* **12** (2023) 001 [arXiv:2210.08382] [INSPIRE].

[127] F. Hernandez Vivanco, R. Smith, E. Thrane, P.D. Lasky, C. Talbot and V. Raymond, *Measuring the neutron star equation of state with gravitational waves: The first forty binary neutron star merger observations*, *Phys. Rev. D* **100** (2019) 103009 [arXiv:1909.02698] [INSPIRE].

[128] J. McGinn, A. Mukherjee, J. Irwin, C. Messenger, M.J. Williams and I.S. Heng, *Rapid neutron star equation of state inference with Normalising Flows*, arXiv:2403.17462 [INSPIRE].

[129] N.K. Glendenning, *Compact stars: Nuclear physics, particle physics, and general relativity*, in *Astronomy and Astrophysics Library*, Springer, New York, NY, U.S.A. (1997) [DOI:10.1007/978-1-4684-0491-3] [INSPIRE].

[130] Y.K. Gambhir, P. Ring and A. Thimet, *Relativistic mean field theory for finite nuclei*, *Ann. Phys.* **198** (1990) 132 [INSPIRE].

[131] J. Boguta and A.R. Bodmer, *Relativistic Calculation of Nuclear Matter and the Nuclear Surface*, *Nucl. Phys. A* **292** (1977) 413 [INSPIRE].

[132] G. Baym, C.J. Pethick and P. Sutherland, *The Ground state of matter at high densities: Equation of state and stellar models*, *Astrophys. J.* **170** (1971) 299 [INSPIRE].

[133] G. Baym, H.A. Bethe and C.J. Pethick, *Neutron star matter*, *Nucl. Phys. A* **175** (1971) 225 [INSPIRE].

[134] W.H. Press, S.A. Teukolsky, W.T. Vetterling and B.P. Flannery, *Numerical Recipes*, 3rd edition, Cambridge University Press (2007).

[135] C. Modi, C. Margossian, Y. Yao, R. Gower, D. Blei and L. Saul, *Variational Inference with Gaussian Score Matching*, arXiv:2307.07849.

JCAP09(2024)009