


PROCEEDING

Deducing the EOS of dense neutron star matter with machine learning

Delaney Farrell¹  | Pierre Baldi² | Jordan Ott² | Aishik Ghosh^{3,4} | Andrew W. Steiner^{5,6} | Atharva Kavitar⁷ | Lee Lindblom⁸ | Daniel Whiteson³ | Fridolin Weber^{1,8}

¹Department of Physics, San Diego State University, San Diego, California, USA

²Department of Computer Science, University of California, Irvine, California, USA

³Department of Physics and Astronomy, University of California, Irvine, California, USA

⁴Physics Division, Lawrence Berkeley National Laboratory, Berkeley, California, USA

⁵Department of Physics and Astronomy, University of Tennessee, Knoxville, Tennessee, USA

⁶Physics Division, Oak Ridge National Laboratory, Tennessee, USA

⁷Department of Computer Science, TU Kaiserslautern, Germany

⁸Center for Astrophysics and Space Sciences, University of California, San Diego, California, USA

Correspondence

Delaney Farrell, Department of Physics, San Diego State University, California, USA.

Email: dfarrell@sdsu.edu

Funding information

National Science Foundation, Grant/Award Numbers: 2012857, AST 19-09490, PHY 21-16686, PHY-2012152; U.S. Department of Energy

Abstract

The interior of a neutron star is a unique astrophysical laboratory for studying matter at extreme densities and pressures beyond what is replicable in terrestrial experiments. While there is no direct way to simulate the interior of these stars, one promising avenue to learning more about the equation of state (EOS) of such matter is through X-rays emitted from the star's surface. The current state-of-the-art method for inference of EOS from a star's X-ray spectra uses piece-wise, simulation-based likelihoods that rely on theoretical assumptions complicated by systematic uncertainties. To reduce the dimensionality of the problem, this method infers macroscopic properties of the star (mass and radius) from emitted X-ray spectra, and from those quantities infers the EOS. This work approaches the same problem using machine learning techniques, demonstrating a series of enhancements to the current state-of-the-art by realistic uncertainty quantification and reducing the need for theoretical assumptions. We also demonstrate novel inference of the EOS directly from high-dimensional simulated X-ray spectra from neutron stars that negate the need for a piece-wise approach. This inference allows for a natural propagation of uncertainties from the X-ray spectra by conditioning the discussed networks on realistic sources of uncertainty for each star.

KEYWORDS

machine learning, neutron stars, nuclear equation of state, X-ray spectra, XSPEC

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs License](https://creativecommons.org/licenses/by-nc-nd/4.0/), which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2023 The Authors. *Astronomische Nachrichten* published by Wiley-VCH GmbH.

1 | INTRODUCTION

Neutron stars are the neutron-packed remnants of massive stars ($\gtrsim 8M_{\odot}$) created in supernova explosions. These stars provide a unique laboratory for studying matter in physical conditions that cannot yet be replicated on Earth: relatively cold, isospin-asymmetric matter at supranuclear densities. Insights about the different forms of matter which emerge under such extreme conditions can improve our understanding of the physics of the early universe, fundamental forces like quantum chromodynamics and gravity, and other astrophysical phenomena.

The nature of matter within a neutron star is summarized by its equation of state (EOS), or the relationship between pressure P and energy density ϵ . Different hypotheses on stable states of exotic matter potentially existing in the inner cores of neutron stars result in starkly different relationships for the nuclear EOS. The EOS of a star in static gravitational equilibrium determines macroscopic stellar properties such as its mass and radius, which influences observables such as the stellar X-ray spectrum. Conversely, the observed stellar spectra can be used to infer macroscopic properties, which in principle allow inference of the EOS (Heinke et al. 2006; Lattimer & Prakash 2001; Lindblom 2010; Lindblom & Indik 2014; Rutledge et al. 1999; Steiner et al. 2010). The inference of EOS from mass and radius is numerically very difficult, as the theoretical framework that relates EOS to mass and radius, the relativistic stellar structure equations, are not easily inverted. The small number of neutron star observations and the significant uncertainty of individual measurements add additional challenges to this problem. It is therefore vital to extract as much information as possible from each observation to provide the best approximation of the star's EOS (Farrell et al. 2023).

We use neural network regression to infer the EOS of neutron star matter. We begin by constructing a piece-wise network following inference used in the state-of-the-art, where EOS is inferred from stellar mass and radius information derived from realistic simulated X-ray spectra. This involves two networks: the first performs inference of stellar masses and radii from high-dimensional X-ray spectra, and the second infers the star's EOS from stellar mass and radius information while accounting for the associated underlying uncertainties. Finally, we perform a novel inference of EOS directly from a multi-star *set* of stellar spectra without the intermediate step of collapsing the spectra information into mass and radius. Each method allows for a full propagation of uncertainties by conditioning the networks on the stellar nuisance parameters (NPs) representative of realistic uncertainties. When possible, comparisons are made to benchmark methods.

2 | BACKGROUND

2.1 | Equation of state of neutron star matter

The relativistic stellar structure equations provide a mathematical framework linking macroscopic characteristics like gravitational mass M and radius R of a neutron star to the stellar matter's EOS. Assuming the object is spherically symmetric, non-rotating, and non-magnetic, these equations take the form of the well-known Tolman–Oppenheimer–Volkoff (TOV) equation. Given an EOS, numerically solving the TOV equation for gravitational mass M and radius R is straightforward. Inverting the equation can provide constraints on the EOS from observable properties, but would require the input of at least two exact mass and radius measurements—something not possible with the current observational technology (Lindblom & Indik 2014). Therefore numerically solving the inverse problem is much more complicated, potentially even intractable without making significant numerical assumptions (Farrell et al. 2023).

2.2 | X-ray spectroscopy

Observation of neutron star emission is perhaps the most promising avenue to learn more about the interior conditions of the star through measurements of mass and radius. Much of the neutron star observational data comes from X-ray emission, either from electromagnetic radiation from pulsars or thermal emission of quiescent low-mass X-ray binaries (qLMXBs) (Farrell et al. 2023). Low-mass binaries historically have provided strong constraints on neutron star structure, as their low magnetic fields result in minimal effects on the temperature distribution and radiation transport on the stars' surface (Bogdanov et al. 2016; Campana et al. 1998; Potekhin 2014). These binaries are also identified in globular clusters where distances, ages, and reddening are well-constrained (Heinke et al. 2003), removing some of the uncertainties in observational measurement. Inference of mass and radius from X-ray spectra is done through the process of spectral fitting, where the emitted X-ray spectrum is fitted to theoretical atmosphere models in the software XSPEC (Arnaud 1996). In this work, the simulated X-ray spectra used in EOS inference are derived from the same fitting software.

3 | TRAINING SAMPLES

Described below are the samples of simulated neutron star data used to train the networks and evaluate their

performance. Each star is described by five quantities: two of which (mass and radius) are drawn from mass-radius relations determined by various EOS models, and the other three are NPs representative of observational uncertainties. The three NPs are independent of the EOS and can vary from star to star. The five stellar parameters are used to determine the simulated X-ray telescope spectrum in the chosen NS theoretical model (more information below in 3.2 and (Farrell et al. 2023)). In the case of EOS inference, sets of stars with consistent EOS are grouped to form training and testing sets.

3.1 | Equation of state

For the EOS of the hadronic model within the core of a neutron star, we begin with the relativistic mean field model GM1L (Typel et al. 2010). The parameterization used in this work accounts only for protons and neutrons but can be extended to include more exotic particles like hyperons and Δ baryons (Spinella & Weber 2020).

A tabulated EOS is unideal as the outcome of regression, so in order to limit the number of parameters each network must learn, the essential features of the high-density portion of the EOS were represented by constructing parametric representations based on spectral fits (Lindblom 2010). A full description of parametric representations of the EOS can be found at (Lindblom 2018). We constructed spectral fits for the GM1L EOS using two parameters, hereafter referred to as λ_1 and λ_2 . The original GM1L parameters were randomly shifted to generate 10^4 different EOS variations (see (Farrell et al. 2023) for more details). Each EOS variation was then used to generate a $M - R$ relation using the TOV equation, resulting in 10^6 (M, R) pairs, each representing stellar parameters consistent with that EOS.

3.2 | Modeling X-ray spectra

The XSPEC program (Arnaud 1996), which is traditionally used spectral fitting, is also capable of generation of simulated X-ray spectra, via the `fakeit` command. For each simulated spectra, the user must supply a theoretical model and a telescope response matrix. For spectra generation, the NS theoretical model NSATMOS (Heinke et al. 2006) was used, which requires five input parameters discussed below. The Chandra telescope response specified by (Heinke et al. 2006) was also used to describe the instrument response and telescope effective area.

TABLE 1 Description of “true”, “tight”, and “loose” NP scenarios.

Nuisance Parameter	True	Loose
Distance, d	exact	20%
Hydrogen Column, N_H	exact	50%
$\log T_{\text{eff}}$	exact	± 0.2

Note: Shown are the width of each Gaussian distribution representing the prior knowledge of each NP. For distance and N_H , width is relative; for $\log(T_{\text{eff}})$, it is absolute. Table adapted courtesy of Journal of Cosmology and Astroparticle Physics (JCAP) (Farrell et al. 2023)

3.3 | Nuisance parameters

The NSATMOS model has five parameters to describe each star: gravitational mass M in units of M_\odot , radius R in units of km, and three additional parameters related to observation. For the context of $M - R$ and subsequent EOS inference, only M and R are parameters of interest, whose values come from those generated by the GM1L EOS and so provide information relevant to the physical question. The remaining three parameters are the effective temperature of the surface, T_{eff} , the distance to the star, d , and the hydrogen column, N_H . These values are a leading source of uncertainty in the inference of mass and radius, and subsequently the EOS, in actual observation but not targets of the regression, so they are referred to as NPs.

We drew values for each NP from realistic ranges seen in observation (see (Farrell et al. 2023) for more details). To demonstrate the impact of NP uncertainties in regressed parameters, we define three example scenarios of uncertainties (“true”, “tight”, and “loose”) which describe the quality of prior information on the NP values for each star. These prior ranges are shown in Table 1, taken from (Farrell et al. 2023). Results for all three NP cases are shown in (Farrell et al. 2023), but for the context of this proceeding, only results for “true” and “loose” are described.

4 | PIECE-WISE APPROACH: INFERENCE OF THE EOS FROM REGRESSED MASS AND RADIUS

The following networks infer EOS parameters by first inferring mass-radius pairs from stellar spectra. We compare neural network inference of mass and radius to the current state of the art, XSPEC.

4.1 | Inference of mass and radius from spectra

Direct analysis of X-ray spectra by neural networks in a novel concept, made complicated by the spec-

tra's high dimensionality, $\mathcal{O}(10^3)$ bins of photon energy. Summarizing the spectrum information into two values, mass and radius, is one avenue to reducing the dimensionality of the EOS inference problem. Following the current state-of-the-art piecewise approach, we begin with constructing a network to estimate the mass and radius from a single stellar spectrum.

4.1.1 | Mass and radius inference with the MR_Net method

We construct a neural network that takes energy bins of the stellar X-ray spectrum as an input and outputs estimates of the star's mass and radius. In addition, the mass and radius regressor is parameterized on the NPs (distance, N_H , $\log T_{\text{eff}}$), which allows the results to be conditioned on the NPs. A full description of the architecture of MR_Net can be found in (Farrell et al. 2023). The performance of mass and radius inference is compared to spectral fitting performed by XSPEC.

4.1.2 | Mass and radius inference by XSPEC

The XSPEC performs mass and radius inference given a sample observed X-ray spectrum by scanning $M - R$ parameter space for the best fitting values. For each mass-radius pair, XSPEC calculates an expected spectrum using the chosen model and telescope response function, identical to those used to generate the sample spectra being fit (Farrell et al. 2023). The best-fit mass-radius pair minimizes a bin-wise χ^2 .

For both XSPEC and MR_Net, inference of mass and radius from a given spectrum is repeated several times with varying assumed values of the NPs drawn from “true” or “loose” priors in order to reflect the lack of knowledge of the NP values in real observation. Performance of MR_Net with comparison to XSPEC is shown below.

4.2 | MR_Net performance in mass, radius

MR_Net proves capable of extracting the mass and radius values of the star directly from a given X-ray spectrum and NPs. MR_Net achieves high performance without the knowledge of the theoretical model used to generate the training examples, while XSPEC requires precise specification of the theoretical model for similar results.

Because MR_Net takes NP values as input, we can assess the impact of NP uncertainty testing on values




FIGURE 1 Performance of the MR_Net regression of a neutron star mass from its stellar X-ray spectrum, compared to regression using XSPEC. Figure credit: IOP Publishing LTD and SiSSa Medialab. Reproduced by permission of IOP Publishing. All rights reserved (Farrell et al. 2023)

drawn from different priors. When not known exactly (the “loose” case), the error residuals widen for both MR_Net and XSPEC as expected. Figures 1 and 2 show the mass and radius residuals, respectively, while Table 2 shows the mean and width of each distribution (Farrell et al. 2023).

4.3 | Inference of the EOS from mass and radius

Deep feed-forward neural networks for each NP scenario are trained to regress the EOS parameters λ_1 and λ_2 from a collection of ten stars. Each star is represented by its mass and radius, provided either from MR_Net or XSPEC; a full description of the networks can be found in (Farrell et al. 2023). The performance of each network (shown visually in Figures 3 and 4) is compared to the network which performs inference of EOS directly from X-ray spectra, described below.




FIGURE 2 Performance of the MR_Net regression of a neutron star radius from its stellar X-ray spectrum, compared to regression using XSPEC. Shown is the residual for two scenarios of nuisance parameter uncertainties. Figure credit: IOP Publishing Ltd and Sissa Medialab. Reproduced by permission of IOP Publishing. All rights reserved (Farrell et al. 2023)

TABLE 2 Performance of the regression of neutron star mass and radius for XSPEC as well as our neural network regression, MR_Net which lacks any knowledge of the theoretical model.

Method	Nuisance Parameters	Mass		Radius	
		μ	σ	μ	σ
XSPEC	True	-0.01	0.50	0.23	1.44
MR_Net	True	-0.14	0.93	-0.07	2.80
XSPEC	Loose	0.18	0.86	-0.06	4.32
MR_Net	Loose	0.28	1.29	0.14	4.93

Note: Shown are the mean (μ) and standard deviation (σ) of the residuals under two scenarios of nuisance parameter uncertainties. Table adapted by permission of IOP Publishing (Farrell et al. 2023)

5 | INFERENCE OF THE EOS FROM SPECTRA

In this section, we introduce a single network to regress EOS parameters λ_1 and λ_2 directly from a set of ten




FIGURE 3 Performance of the regression of neutron star EOS parameter λ_1 using direct regression from spectra, as compared to regression from mass and radius information extracted via MR_Net or XSPEC. Shown are the residual distributions, the difference between the true and predicted values, under two scenarios of nuisance parameter uncertainties. See Table 3 for quantitative analysis. Figure credit: IOP Publishing LTD and Sissa Medialab. Reproduced by permission of IOP Publishing. All rights reserved (Farrell et al. 2023)

stellar spectra. The intention of performing the regression directly on X-ray spectra is to keep as much information from observation as possible, while also allowing for robust uncertainty quantification with the propagation of NPs.

Unlike the deep neural networks used in the piece-wise approach, this network adopts the transformer architecture in order to operate on an unordered set of vectors (independent spectra observations for multiple stars). For a full description of the transformer architecture, see (Farrell et al. 2023; Vaswani et al. 2017). This networks takes X-ray spectra and corresponding NPs for ten individual stars, though the size of the input stellar cluster can be modified with minimal effort due to the transformer architecture. The final output of the network is the two EOS parameters for the stellar cluster, λ_1 and λ_2 .




FIGURE 4 Performance of the regression of neutron star EOS parameter λ_2 using direct regression from spectra, as compared to regression from mass and radius information extracted via MR_Net or XSPEC. Figure credit: IOP Publishing Ltd and Medialab. Reproduced by permission of IOP Publishing. All rights reserved (Farrell et al. 2023)

5.1 | Results

Figures 3 and 4 show the residuals in the EOS parameters for the spectra to EOS regression, as compared to regression from mass and radius information provided by MR_Net or XSPEC. Table 3 summarizes the performance for each method.

6 | DISCUSSION

We have demonstrated the use of different neural network architectures to regress EOS parameters from simulated neutron star X-ray spectra. Two routes are taken: the first follows the current state-of-the-art, employing a piece-wise approach inferring a star’s mass and radius from spectra, and using those values to infer EOS parameters for a cluster of ten stars. The second approach directly

TABLE 3 Performance of the regression of neutron star EOS parameters λ_1 and λ_2 using direct regression from spectra, as compared to NN regression from mass and radius (M, R) information extracted via MR_Net or XSPEC.

Method	Nuisance Parameters	λ_1		λ_2	
		μ	σ	μ	σ
NN(Spectra)	True	-0.02	0.066	0.01	0.075
NN(M, R via MR_Net)	True	-0.03	0.089	-0.02	0.068
NN(M, R via XSPEC)	True	-0.03	0.065	0.01	0.055
NN(Spectra)	Loose	-0.03	0.131	-0.01	0.078
NN(M, R via MR_Net)	Loose	-0.01	0.135	-0.02	0.078
NN(M, R via XSPEC)	Loose	-0.03	0.123	0.01	0.058

Note: Shown are the mean (μ) and standard deviation (σ) of the residuals under two scenarios of nuisance parameter uncertainties. Table adapted by permission of IOP Publishing (Farrell et al. 2023).

regresses EOS parameters from a the stellar spectra of 10 stars, a novel approach which directly analyzes sets of high-dimensional telescope spectra. The networks are conditioned on the NPs in each step, folding in NP uncertainty by multiple sampling from priors (Farrell et al. 2023). The results demonstrate machine learning techniques achieve comparable performance to methods that rely on theoretical models used to generate the simulated samples.

ACKNOWLEDGMENTS

DW and AG are supported by The Department of Energy Office of Science. LL was supported by NSF Grant No. 2012857 to the University of California at San Diego. AWS was supported by NSF AST 19-09490, PHY 21-16686, and the Department of Energy Office of Nuclear Physics. DF and FW are supported by the National Science Foundation (USA) under Grant No. PHY-2012152. The authors are grateful to C. O. Heinke and W. C. G. Ho for their assistance in understanding and navigating the XSPEC software.

ORCID

Delaney Farrell  <https://orcid.org/0000-0003-1512-712X>

REFERENCES

- Arnaud, K. A. 1996, in: *Astronomical Data Analysis Software and Systems V*, eds. G. H. Jacoby & J. Barnes, Astronomical Society of the Pacific (San Francisco, CA), Vol. 101, 17.
- Bogdanov, S., Heinke, C. O., Özel, F., & Güver, T. 2016, *ApJ*, 831(2), 184.
- Campana, S., Colpi, M., Mereghetti, S., Stella, L., & Tavani, M. 1998, *Astron. Astrophys. Rev.*, 8(4), 279.

- Farrell, D., Baldi, P., Ott, J. et al. 2023, arXiv:2209.02817, To appear in the Journal of Cosmology and Astroparticle Physics (JCAP).
- Heinke, C. O., Grindlay, J. E., Lugger, P. M., Cohn, H. N., Edmonds, P. D., Lloyd, D. A., & Cool, A. M. 2003, *ApJ*, 598(1), 501.
- Heinke, C. O., Rybicki, G. B., Narayan, R., & Grindlay, J. E. 2006, *ApJ*, 644(2), 1090.
- Lattimer, J. M., & Prakash, M. 2001, *ApJ*, 550, 426.
- Lindblom, L. 2010, *Phys. Rev. D*, 82, 103011.
- Lindblom, L. 2018, *Phys. Rev. D*, 97, 123019.
- Lindblom, L., & Indik, N. M. 2014, *Phys. Rev. D*, 89(6), 064003 ([Erratum: *Phys. Rev. D* 93, 129903 (2016)]).
- Potekhin, A. Y. 2014, *Physics-Uspekhi*, 57(8), 735.
- Rutledge, R. E., Bildsten, L., Brown, E. F., Pavlov, G. G., & Zavlin, V. E. 1999, *ApJ*, 514(2), 945.
- Spinella, W. M., & Weber, F. 2020, in: *Dense Baryonic Matter in the Cores of Neutron Stars*, ed. C. A. Z. Vasconcellos, World Scientific (Singapore), 85.
- Steiner, A. W., Lattimer, J. M., & Brown, E. F. 2010, *ApJ*, 722, 33.
- Typel, S., Röpke, G., Klähn, T., Blaschke, D., & Wolter, H. H. 2010, *Phys. Rev. C*, 81, 015803.

- Vaswani, A., Shazeer, N., Parmar, N., et al. 2017, *Adv. Neural. Inf. Process. Syst.*, 30, 5998.

AUTHOR BIOGRAPHY

Delaney Farrell is a PhD candidate in computational science at San Diego State University and the University of California, Irvine. She received her bachelor's degree in physics at San Diego State University. She is interested in neutron stars, dense matter, and computational astrophysics.

How to cite this article: Farrell, D., Baldi, P., Ott, J., et al. 2023, *Astron. Nachr.*, e20230009. <https://doi.org/10.1002/asna.20230009>