

Searching for phase transitions in neutron stars with modified Gaussian processes

Debora Mroczek¹, M. Coleman Miller², Jacquelyn Noronha-Hostler¹
and Nicolas Yunes¹

¹ Illinois Center for Advanced Studies of the Universe, Department of Physics, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

² Department of Astronomy, University of Maryland, College Park, 20742, MD, USA

E-mail: deboram2@illinois.edu

Abstract. Gaussian processes provide a promising framework by which to extrapolate the equation of state (EoS) of cold, catalyzed matter beyond 1 – 2 times nuclear saturation density. Here we discuss how to extend Gaussian processes to include nontrivial features in the speed of sound, such as bumps, kinks, and plateaus, which are predicted by nuclear models with exotic degrees of freedom. Using a fully Bayesian analysis incorporating measurements from X-ray sources, gravitational wave observations, and perturbative QCD results, we show that these features are compatible with current constraints and report on how the features affect the EoS posteriors.

1. Introduction

Numerous types of electromagnetic and gravitational-wave data sets are now available for neutron stars, and these are being used to infer the composition, thermodynamic, and dynamical properties of the core matter of these stars (see, e.g., Refs. [1, 2, 3, 4, 5, 6]). Upgrades to current gravitational wave detectors as well as future instruments and results from the Neutron star Interior Composition Explorer (NICER) mission are expected to improve the data available and, therefore, tighten the constraints on the equation of state (EoS) of cold, catalyzed nuclear matter in neutron stars.

An important part of the inference procedure, which is still under active debate, is how to generate minimally biased candidate EoS that incorporate all relevant physics. Parametric models such as piecewise polytropes or spectral decompositions are commonly used, but these can introduce unwanted correlations across density scales [7]. Additionally, sharp features in the speed of sound – which are present in nuclear models with exotic degrees of freedom (see, e.g. Ref. [8]) – are not well represented in the most common parameterizations, although they play an important role in understanding heavy, ultra-heavy, and twin stars [9, 10].

Instead, we start with procedures based on Gaussian processes (GP), which are more model-agnostic and can mitigate these issues at the cost of increased functional complexity [11]. It has been shown that GP ensembles capture a wide range of features in the EoS, including phase transitions of varying strength at different densities (see Refs. [12, 13] for detailed discussions). However, investigating specific features in the EoS with GPs can pose computational and sampling challenges.

Here, as a way to produce priors that adequately capture behavior that is predicted by state-of-the-art nuclear models at a low computational cost and without introducing model dependencies, we present a method for generating EoS with sharp features, which we call modified Gaussian processes (mGP). An mGP sample consists of a smooth baseline EoS sampled from a tailored GP that is modified over some range in pressure. The modifications are constructed such that we produce bumps, spikes, oscillations, plateaus, and first-order phase transitions while keeping track of how and where modifications appear. We generate a prior distribution containing samples from both a GP with fixed hyperparameters and a modified GP and construct mass-radius and EoS posteriors using available measurements and theoretical input based on the Bayesian procedure described in Ref. [14]. The different constructions lead to nearly identical mass-radius posteriors, but the EoS posterior is wider in the mGP than in the unmodified GP framework for densities above ~ 1.5 times the nuclear saturation density ($n_{\text{sat}} = 0.15 \pm 0.01 \text{ fm}^{-3}$). In addition, the Bayesian evidence for these non-smooth EoS is indistinguishable from that for smooth EoS. Thus, current measurements are consistent with sharp and non-trivial features in the EoS, but the data is not yet informative enough to favor or disfavor them.

2. Generating the Equation of State

In this section, we give a brief review of GP EoS and describe how we perform our modifications.

Generally, a Gaussian process yields the joint probability density for a continuous function $f(x)$ at domain points $\vec{x} = \{x_i\}$, which is a tool for approximating f over some domain. This joint probability density is presumed to be a multivariate Gaussian with means $\vec{\mu}$ and covariance matrix Σ for the set of domain points \vec{x} . That is,

$$f(\vec{x}) \sim \mathcal{N}(\vec{\mu}, \Sigma), \quad (1)$$

where \mathcal{N} indicates a normal distribution. The means and covariance matrix are specified by hyperparameters and depend on further assumptions about $f(x)$.

In this context, a given functional form for the neutron star matter EoS can be sampled from a GP which approximates the relationship between two thermodynamic quantities. We follow Refs. [15, 11] in introducing the auxiliary variable

$$\phi \equiv \ln \left(\frac{d\epsilon}{dP} - 1 \right), \quad (2)$$

where ϵ is the energy density and P is the pressure. For any $\phi \in (-\infty, \infty)$, the adiabatic speed of sound, $c_s^2 = dP/d\epsilon$, is stable and causal by construction, since when $\phi \rightarrow +\infty$, $c_s^2 \rightarrow 0$, and when $\phi \rightarrow -\infty$, $c_s^2 \rightarrow 1$, in units where the speed of light $c \equiv 1$.

We model ϕ as a function of $\log_{10} P$ in cgs units at 100 equally spaced points in the range $\log_{10} P (\text{erg cm}^{-3}) \sim 32 - 38$, using the following trend for the means

$$\mu_i(\log_{10} P_i) = b - 2(\log_{10} P_i - 32.7), \quad (3)$$

where $b = 5.5$ or 3.7 (sampled with equal probability). The first value of b corresponds to the approximate trend found in Ref. [5] using several tabulated nuclear EoS. This trend largely produces EoS that approach the causal limit in the range of maximal central densities. The second value of b produces EoS that tend to much lower values of c_s^2 and are near the conformal limit $c_s^2 \rightarrow 1/3$ at central densities for maximally massive stars.

We assume that the correlation between the function values at different domain points can be represented by a kernel function of the two points, $\Sigma(f_i, f_j) = K(x_i, x_j)$. We also assume that

$K(x_i, x_j)$ is a Gaussian which depends only on the distance between the two points, commonly referred to as the squared-exponential kernel,

$$K_{\text{se}}(x_i, x_j) = \sigma^2 \exp \left[-\frac{(x_i - x_j)^2}{2l^2} \right], \quad (4)$$

where l determines the correlation length scale (the limit where all points are independent of each other is $l \rightarrow 0$) and σ is the strength of the overall correlation ($\sigma \rightarrow 0$ means the variance is negligible). We choose to fix $l = \sigma = 1$, following Ref. [5], which is motivated by the spread of the tabulated EoS used to obtain Eq. (3). Our approach differs from that in Ref. [5] in that we allow for two values of b , which introduces an additional degree of flexibility and ensures that a wider range of c_s^2 are represented, despite the other hyperparameters being fixed.

A modified Gaussian process EoS is produced by taking a sample from the GPs outlined above, which becomes the baseline EoS, removing *part* of the baseline EoS over some range in P , and then replacing it with a modification. Modifications are introduced in the form of spikes and plateaus in three different categories:

- (i) Sharp bump/dip – a single point in the EoS (selected randomly with uniform probability) gets shifted up/down with respect to the baseline to a random value of c_s^2 between 0 and 1 that is at least 10% above or below the baseline value.
- (ii) Bump/dip + plateau – In addition to a bump/dip, a plateau is introduced. A plateau spans between 2 and 20 points in pressure (in units of erg cm^{-3} , $\log_{10} \Delta P_{\text{plateau}} = 0.12 - 1.2$). The location, height, and length of a plateau are sampled randomly with uniform probability across the allowed ranges in P and c_s^2 . There is a conditional relationship between the bump/dip and the plateau: if a bump above the baseline value is introduced, the plateau must be below the baseline value at the point where the plateau is introduced. If a dip is introduced, the plateau must be above the baseline value. The 10% rule of the sharp bump/dip modification still applies.
- (iii) Double plateau – In this case two plateaus are introduced, instead of a bump/dip + plateau. The same rules apply.

We produce a total of 900,000 candidate EoS. For any given sample there is a 25% probability that it will be left unmodified and a 75% probability that it will be modified in one of the three ways listed above (25% for each category). At densities below $0.5n_{\text{sat}}$ (roughly the crust-core transition density [16]), we use the QHC19 EoS [17]. We also compute the full range of thermodynamic quantities and remove any mGP samples that contain modifications below the nuclear saturation density, where the properties of equilibrium nuclear matter are better constrained.

Figure 1 illustrates samples from both the Gaussian process and the modified Gaussian process, where the speed of sound is shown as a function of the pressure. Through this small sample of EoS, we see that the modification procedure introduces various interesting features consistent with nuclear models with exotic degrees of freedom, at a low computational cost. We emphasize that traditional GPs can, in principle, approximate these functions and that the unmodified GPs implemented here are tailored to two specific trends and are, thus, not representative of the capabilities of GPs in general. However, the low computational cost and ability to track the location and properties of nontrivial features associated with mGPs make this framework useful for studying exotic degrees of freedom in neutron star cores.

3. Prior pruning based on observational constraints

Before proceeding with the Bayesian analysis, we perform a process of informing our prior with conservative cuts motivated by NICER measurements of PSR J0030+0451 [18] and




Figure 1. The speed of sound squared in units of c^2 , as a function of the pressure, for an unmodified (left panel) and modified (right panel) Gaussian process. We use a squared-exponential kernel with fixed hyperparameters, $\sigma = 1$, $l = 1$, and $\mu(\log_{10} P) = b - 2(\log_{10} P - 32.7)$, where P is the pressure in cgs units and $b = 5.5$ or 3.7 (see text for more details). These samples have not been compared with data, which means that the EoS need not be consistent with observational and theoretical constraints.

J0740+6620[5], and the detection of GW170817 [19]. We refer to this step as *pruning*, as it is used not to extract the EoS, but rather to ensure we have a suitable number of strong candidate EoS in our prior.

The pruning process happens in two steps. We first impose that the maximum mass of a nonrotating star, M_{\max} , must be at least $1.8 M_{\odot}$. For these EoS, we then calculate the equatorial circumferential radius and dimensionless tidal deformability for a $1.4 M_{\odot}$ star, $R_{1.4}$ and $\Lambda_{1.4}$, respectively. We require that $10 \leq \Lambda_{1.4} \leq 2000$ and $9.0 \leq R_{1.4} \leq 18.0$ km.

In Figure 2, we show the effect of the prior pruning on the speed of sound of the combined GP and mGP samples. The priors are labeled 1 and 2 and are represented by their means and 90% contours. Prior 1 is the original sample of 900,000 EoS (25/75 % split between GP/mGP samples) which does not incorporate assumptions about the physics. Prior 2 shows the 104,594 samples (30/70 % split between GP/mGP samples) which meet the requirements informed by GW190817 and PSR J0030+0451, the mass cut $M_{\max} \geq 1.8 M_{\odot}$, and had no modifications below n_{sat} . This final set of EoS is overall stiffer compared to prior 1 but constrained to low values of c_s^2 for $\log_{10} P \lesssim 33.5$ because of the constraints at $1.4 M_{\odot}$ and theoretical input. We use prior 2 in the Bayesian analysis.

With Figure 2, we also note that implementing some form of prior pruning is important in nonparametric frameworks for the EoS, especially in the model-agnostic limit. Because these frameworks incorporate minimal physics by construction, creating a reasonable subset of EoS takes significantly more samples. Taking the numbers from our analysis, out of an ensemble of 900,000 EoS, only $\sim 10\%$ provide a reasonable match to the observed properties of neutron stars.

4. Statistical method and constraints

In the Bayesian analysis, we follow the procedure outlined in detail in Ref. [5]. As usual, the posterior probability of a given EoS model k is proportional to the product of the prior probability and the likelihood, $p_k \propto q_k \mathcal{L}_k$. The likelihood of a set of observations (i, j) for EoS model k is given by




Figure 2. The speed of sound squared in units of c^2 as a function of the pressure from the combined distribution of GP and mGP samples. The solid lines delimit the 5% and 95% contours and the dashed lines indicate the mean. Prior 1 (gray) corresponds to 900,000 candidate EoS before any pruning is performed. Prior 2 (blue) contains the 104,594 EoS that meet the conservative criteria for the mass and tidal deformability based on GW170817 and PSR J0030+0451 and $M_{\max} \geq 1.8M_{\odot}$.

$$\mathcal{L}_k = \prod_i \left[\prod_{j=1}^{j(i)} \mathcal{L}_k(i, j) \right] \quad (5)$$

where i is a type of measurement (e.g. mass, radius) and j is an independent measurement of type i (e.g. two independent measurements of the mass of one object). We follow the procedure developed in Ref. [14] to determine \mathcal{L}_k for each EoS, which uses the full posterior probability distributions of all measurements included.

The constraints included in the analysis are the masses of three high-mass pulsars [20, 21, 22], the joint mass and tidal deformability posteriors extracted from GW170817 [19] and GW190425 [23], and the mass-radius posterior from NICER data on PSR J0030-0451 [18] and PSR J0740+6620 [5] (see, respectively, [24] and [25] for independent analyses of these two pulsars from a separate group within the NICER collaboration), with a Gaussian prior for the value of the symmetry energy at nuclear saturation density, $S = 32 \pm 2$ MeV [26]. We also implement perturbative QCD (pQCD) stability and integral constraints from Ref. [27] for the renormalization parameter range $X = [1/2, 2]$, using the pQCD likelihood introduced in Ref. [2] at the central density for a maximally massive star, n_B^{max} (varies for each EoS), as was done in [28].

5. Results

We separate the EoS in the prior between modified and unmodified samples and look at the Bayesian evidence (the probability for a model given the constraints) for each model separately.

Since we assume each EoS has an equal prior probability, the Bayesian evidence for a smooth EoS (represented by the GP samples) or an EoS with sharp/non-trivial features (mGP samples) is given by the average likelihood of the EoS in each category,

$$P[m(GP)|\text{constraints}] = \frac{1}{N_{m(GP)}} \sum_{i=0}^{N_{(m)GP}} \mathcal{L}_i, \quad (6)$$

where the prior probability of a model is $P[m(GP)] = \frac{1}{N_{m(GP)}}$ and \mathcal{L}_i is the likelihood of EoS i .

The Bayes factor K is the ratio between the evidence for hypotheses – the bigger the deviation from unity, the more indication that a set of data supports one hypothesis over the other. Comparing the evidence for GP EoS against mGP EoS yields

$$K = \frac{P[(GP)|\text{constraints}]}{P[(mGP)|\text{constraints}]} = 1.126, \quad (7)$$

from which we conclude that sharp and non-trivial features in the EoS are not ruled out by current constraints. Note also that $K=1.126$ is a statistically insignificant deviation from unity. That is, each hypothesis is comparably good at explaining the data.

In Figure 3, we show the mass-radius prior and posterior probability distributions for GP and mGP EoS separately and find that they are virtually the same. Rather than approximating a continuous distribution from the samples, we bin the EoS by mass and radius and then normalize such that the sum of the heights of the bins is equal to one. To produce the posterior, we weigh each EoS by its corresponding likelihood before normalizing the bins. We also show the 90% and 68% credible regions for the radius at a given mass for $1.1 \leq M \leq 3.5 M_\odot$. Note that, although we show the credible regions for masses as high as $3.5 M_\odot$, the posterior probability is negligible for masses $M \gtrsim 2.7 M_\odot$. We emphasize that the radius contours shift to the right for large masses because larger maximum masses require stiffer EoS, and thus larger radii. Hence, the shape of the contours is not representative of individual EoS, because the only EoS that contribute in the high-mass region are the ones that reach high masses.

In Figure 4, we show the EoS posteriors for the square of the speed of sound as a function of baryon number density in units of n_{sat} , up to n_B^{max} for each EoS. The binning and calculation of the prior and posterior probability distributions are done as discussed for Figure 3. Also shown are the 68% and 90% credible regions for the speed of sound at a given density up to $n_B = 8.0 n_{\text{sat}}$, although the posterior probability for $n_B^{\text{max}} \gtrsim 6.0 n_{\text{sat}}$ is negligible. As before, the only EoS that contribute at any given density are the ones that are still stable neutron stars at that particular density.

Overall, the posteriors are in good agreement. We note that, at the 90% level, compared to the GP posterior, the mGP posterior is wider. Notably, the mGP posterior allows for slightly stiffer EoS in the regime $1.5 \lesssim n_B \lesssim 3.0 n_{\text{sat}}$, and softer EoS above $3.0 n_{\text{sat}}$. For instance, at twice nuclear saturation density, we extract $c_s^2 = 0.29_{-0.11}^{+0.27}$ using EoS from the GP and $c_s^2 = 0.29_{-0.14}^{+0.34}$ using the mGP posterior. At four times nuclear saturation density, the GP EoS range is $c_s^2 = 0.63_{-0.23}^{+0.27}$, while the mGP EoS allows for $c_s^2 = 0.59_{-0.34}^{+0.31}$. While the relative differences in the posteriors are dependent on our choices for the two GPs implemented in the unmodified case, the point holds that when the speed of sound is allowed to display sharp features, the posteriors are wider than when a smooth EoS is presumed. We, therefore, argue that such features should be adequately represented in priors for the extraction of the EoS in neutron star regimes. We emphasize that adequate representation requires not only that these features are present in the synthetic EoS, but that additional steps are taken to ensure a sufficient number of EoS with nontrivial features can account for observations (see discussion in Section 3 regarding prior pruning).




Figure 3. Mass-radius prior (left) and posterior (right) probability distributions for GP (top) and mGP (bottom) EoS. Both the prior and posterior probability distributions are produced by binning the EoS by mass and radius and then normalizing the heights of the bins such that their sum is equal to one. For the posteriors, each EoS is weighted by the corresponding likelihood. Also shown in the posterior plots are the 90% and 68% credible regions for the radius at a given mass for $1.1 \leq M \leq 3.5 M_{\odot}$.

6. Discussion

We introduced modified Gaussian processes as a new technique for generating synthetic equations of state for the cold, catalyzed nuclear matter in neutron stars. These EoS can be generated at a low computational cost, and contain non-trivial structure consistent with the emergence of exotic degrees of freedom. We prune our samples to ensure that a prior of synthetic EoS contains a large number of strong candidates, given that model-agnostic methods contain minimal physical input by construction. In this analysis, only $\sim 10\%$ of the EoS generated in a combined ensemble of modified and unmodified Gaussian process samples were candidates for explaining current mass, radius, and tidal deformability measurements, reinforcing the importance of informing the prior.

We implement the modified Gaussian process framework in a fully Bayesian analysis to verify




Figure 4. EoS prior (left) and posterior (right) probability distributions for GP (top) and mGP (bottom) EoS. The EoS are represented by the speed of sound squared in units of c^2 as a function of baryon number density in units of n_{sat} . The prior and posterior probability distributions are produced by binning the EoS by the speed of sound and number density and then normalizing the heights of the bins such that their sum is equal to one. For the posteriors, each EoS is weighed by the corresponding likelihood. Also shown in the posterior plots are the 90% and 68% credible regions for the speed of sound squared at a given density for $0.5 \leq n_B \leq 8.0 n_{\text{sat}}$. The posterior probability that the central density for a maximally massive star is greater than $\sim 6.0 n_{\text{sat}}$ is negligible in both cases.

that sharp and non-trivial features are consistent with current measurements and theoretical constraints by comparing the evidence for modified Gaussian process EoS against smooth EoS from a standard Gaussian process. Additionally, at the 90% level, the inclusion of non-trivial structure broadens the allowed range for the speed of sound in the core of neutron stars compared to the assumption that the EoS is smooth.

The results presented here suggest that current constraints are not enough to rule definitively in favor of or against phase transitions to exotic degrees of freedom in the core of neutron stars. The question of how future measurements can further constrain the nuclear EoS has been explored (see, e.g., Refs. [29, 30]), but unambiguous signatures of structure in the EoS are still being investigated. Recently, the slope of binary Love relations has been shown to encode structure in the EoS below $3 n_{\text{sat}}$, with twin stars leading to characteristic signatures that can be detected with future instruments [31]. However, a more complete picture of the phase structure of QCD for asymmetric matter at zero temperature will require constraints from both astrophysical observations and laboratory measurements, as well as input from effective theories and pQCD (for a detailed discussions see Refs. [32, 33]).

Acknowledgments

D.M is supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE – 1746047, the Illinois Center for Advanced Studies of the Universe Graduate Fellowship, and the University of Illinois Graduate College Distinguished Fellowship. J.N.H, D.M., and N.Y. were supported in part by the National Science Foundation (NSF) within the framework of the MUSES collaboration, under grant number OAC2103680. J.N.H. acknowledges financial support from the US-DOE Nuclear Science Grant No. DESC0020633. M.C.M. was supported in part by NASA ADAP grant 80NSSC21K0649. N.Y. was supported in part by NSF Award PHYS-2207650. The authors also acknowledge support from the Illinois Campus Cluster, a computing resource that is operated by the Illinois Campus Cluster Program (ICCP) in conjunction with the National Center for Supercomputing Applications (NCSA), which is supported by funds from the University of Illinois at Urbana-Champaign.

References

- [1] Peter T. H. Pang et al. NMMA: A nuclear-physics and multi-messenger astrophysics framework to analyze binary neutron star mergers. 5 2022.
- [2] Tyler Gorda, Oleg Komoltsev, and Aleksi Kurkela. Ab-initio QCD calculations impact the inference of the neutron-star-matter equation of state. 4 2022.
- [3] Ming-Zhe Han, Jin-Liang Jiang, Shao-Peng Tang, and Yi-Zhong Fan. Bayesian Nonparametric Inference of the Neutron Star Equation of State via a Neural Network. *Astrophys. J.*, 919(1):11, 2021.
- [4] G. Raaijmakers, S. K. Greif, K. Hebeler, T. Hinderer, S. Nissanke, A. Schwenk, T. E. Riley, A. L. Watts, J. M. Lattimer, and W. C. G. Ho. 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(2):L29, 2021.
- [5] M. C. Miller et al. The Radius of PSR J0740+6620 from NICER and XMM-Newton Data. *Astrophys. J. Lett.*, 918(2):L28, 2021.
- [6] Rahul Somasundaram, Ingo Tews, and Jérôme Margueron. Investigating signatures of phase transitions in neutron-star cores. *Phys. Rev. C*, 107(2):025801, 2023.
- [7] Isaac Legred, Katerina Chatziioannou, Reed Essick, and Philippe Landry. Implicit correlations within phenomenological parametric models of the neutron star equation of state. *Phys. Rev. D*, 105(4):043016, 2022.
- [8] A. Clevinger, J. Corkish, K. Aryal, and V. Dexheimer. Hybrid equations of state for neutron stars with hyperons and deltas. *Eur. Phys. J. A*, 58(5):96, 2022.
- [9] Hung Tan, Jacquelyn Noronha-Hostler, and Nico Yunes. Neutron Star Equation of State in light of GW190814. *Phys. Rev. Lett.*, 125(26):261104, 2020.
- [10] Hung Tan, Travis Dore, Veronica Dexheimer, Jacquelyn Noronha-Hostler, and Nicolás Yunes. Extreme matter meets extreme gravity: Ultraheavy neutron stars with phase transitions. *Phys. Rev. D*, 105(2):023018, 2022.

- [11] Philippe Landry and Reed Essick. Nonparametric inference of the neutron star equation of state from gravitational wave observations. *Phys. Rev. D*, 99(8):084049, 2019.
- [12] Reed Essick, Philippe Landry, and Daniel E. Holz. Nonparametric Inference of Neutron Star Composition, Equation of State, and Maximum Mass with GW170817. *Phys. Rev. D*, 101(6):063007, 2020.
- [13] Isaac Legred, Katerina Chatziioannou, Reed Essick, Sophia Han, and Philippe Landry. Impact of the PSR J0740+6620 radius constraint on the properties of high-density matter. *Phys. Rev. D*, 104(6):063003, 2021.
- [14] M. Coleman Miller, Cecilia Chirenti, and Frederick K. Lamb. Constraining the equation of state of high-density cold matter using nuclear and astronomical measurements. 4 2019.
- [15] Lee Lindblom. Spectral Representations of Neutron-Star Equations of State. *Phys. Rev. D*, 82:103011, 2010.
- [16] K. Hebeler, J. M. Lattimer, C. J. Pethick, and A. Schwenk. Equation of state and neutron star properties constrained by nuclear physics and observation. *Astrophys. J.*, 773:11, 2013.
- [17] Gordon Baym, Shun Furusawa, Tetsuo Hatsuda, Toru Kojo, and Hajime Togashi. New Neutron Star Equation of State with Quark-Hadron Crossover. *Astrophys. J.*, 885:42, 2019.
- [18] 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(1):L24, 2019.
- [19] Soumi De, Daniel Finstad, James M Lattimer, Duncan A Brown, Edo Berger, and Christopher M Bower. Tidal deformabilities and radii of neutron stars from the observation of gw170817. *Physical review letters*, 121(9):091102, 2018.
- [20] Zaven Arzoumanian, PT Baker, Adam Brazier, Sarah Burke-Spolaor, SJ Chamberlin, Shami Chatterjee, Brian Christy, James M Cordes, Neil J Cornish, Fronefield Crawford, et al. The nanograv 11 year data set: pulsar-timing constraints on the stochastic gravitational-wave background. *The Astrophysical Journal*, 859(1):47, 2018.
- [21] John Antoniadis, Paulo CC Freire, Norbert Wex, Thomas M Tauris, Ryan S Lynch, Marten H Van Kerkwijk, Michael Kramer, Cees Bassa, Vik S Dhillon, Thomas Driebe, et al. A massive pulsar in a compact relativistic binary. *Science*, 340(6131):1233232, 2013.
- [22] H Thankful Cromartie, Emmanuel Fonseca, Scott M Ransom, Paul B Demorest, Zaven Arzoumanian, Harsha Blumer, Paul R Brook, Megan E DeCesar, Timothy Dolch, Justin A Ellis, et al. Relativistic Shapiro delay measurements of an extremely massive millisecond pulsar. *Nature Astronomy*, 4(1):72–76, 2020.
- [23] B. P. Abbott et al. GW190425: Observation of a Compact Binary Coalescence with Total Mass $\sim 3.4M_{\odot}$. *Astrophys. J. Lett.*, 892(1):L3, 2020.
- [24] Thomas E. Riley et al. A *NICER* View of PSR J0030+0451: Millisecond Pulsar Parameter Estimation. *Astrophys. J. Lett.*, 887(1):L21, 2019.
- [25] Thomas 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(2):L27, 2021.
- [26] M. B. Tsang et al. Constraints on the symmetry energy and neutron skins from experiments and theory. *Phys. Rev. C*, 86:015803, 2012.
- [27] Oleg Komoltsev and Aleksii Kurkela. How Perturbative QCD Constrains the Equation of State at Neutron-Star Densities. *Phys. Rev. Lett.*, 128(20):202701, 2022.
- [28] Rahul Somasundaram, Ingo Tews, and Jérôme Margueron. Perturbative QCD and the Neutron Star Equation of State. 4 2022.
- [29] Philippe Landry, Reed Essick, and Katerina Chatziioannou. Nonparametric constraints on neutron star matter with existing and upcoming gravitational wave and pulsar observations. *Phys. Rev. D*, 101(12):123007, 2020.
- [30] Christian Ecker and Luciano Rezzolla. Impact of large-mass constraints on the properties of neutron stars. *Mon. Not. Roy. Astron. Soc.*, 519(2):2615–2622, 2022.
- [31] Hung Tan, Veronica Dexheimer, Jacquelyn Noronha-Hostler, and Nicolas Yunes. Finding Structure in the Speed of Sound of Supranuclear Matter from Binary Love Relations. *Phys. Rev. Lett.*, 128(16):161101, 2022.
- [32] Alessandro Lovato et al. Long Range Plan: Dense matter theory for heavy-ion collisions and neutron stars. 11 2022.
- [33] Agnieszka Sorensen et al. Dense Nuclear Matter Equation of State from Heavy-Ion Collisions. 1 2023.