


INFERENCE OF THE NEUTRON STAR MATTER EQUATION OF STATE

Len Jonah Tristan Brandes
July, 2024



INFERENCE OF THE NEUTRON STAR MATTER EQUATION OF STATE

Len Jonah Tristan Brandes

Vollständiger Abdruck der von der TUM School of Natural Sciences der Technischen Universität München zur Erlangung des akademischen Grades eines

Doktors der Naturwissenschaften

genehmigten Dissertation.

Vorsitz:

Prof. Dr. Laura Fabbietti

Prüfende der Dissertation:

1. apl. Prof. Dr. Norbert Kaiser
2. Prof. Dr. Lukas Heinrich
3. Prof. Dr. Achim Schwenk

Die Dissertation wurde am 15.07.2024 bei der Technischen Universität München eingereicht und durch die TUM School of Natural Sciences am 10.09.2024 angenommen.



This work has been 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. The author thanks the Japan Society for the Promotion of Science (JSPS) for their support during his two International Fellowships for Research in Japan in the summer of 2023 and in the spring of 2024.

ABSTRACT

Information on the phase structure of strongly interacting matter at high baryon densities can be gained from observations of neutron stars and their detailed analysis. In the present work, Bayesian inference methods are used to set constraints on the speed of sound and hence the equation of state of neutron star matter based on recent multimessenger data in combination with limiting conditions from theoretical calculations. In view of the strong constraints imposed by the heaviest known pulsars, the speed of sound must increase rapidly to ensure the stability of these extreme objects. This required stiffness limits the possible appearance of strong phase transitions in neutron star cores, which are characterized by low sound speeds and extended phase coexistence regions. Given these data-based constraints, much discussed issues such as the quest for a phase transition towards restored chiral symmetry and the active degrees of freedom in cold and dense baryonic matter are reexamined. Furthermore, we develop a novel methodology that can infer the equation of state directly from detector observations without a previously necessary intermediate step. This method relies on the use of neural simulation-based inference methods to learn the analytically unavailable likelihood of the neutron star data. Our novel approach reduces the numerical cost of including additional observations, allowing the inference to scale with the anticipated surge of neutron star data in the coming years.

ZUSAMMENFASSUNG

Informationen über die Phasenstruktur von stark wechselwirkender Materie bei hohen Baryondichten können aus Neutronensternbeobachtungen und deren detaillierter Analyse gewonnen werden. In dieser Arbeit wird die Schallgeschwindigkeit und somit die Zustandsgleichung von Neutronensternmaterie durch Bayes'sche Inferenzmethoden auf der Basis aktueller Multimessenger-Daten in Kombination mit limitierenden Ergebnissen aus theoretischen Berechnungen eingeschränkt. Für die Stabilität der schwersten bekannten Pulsare ist ein rapider Anstieg der Schallgeschwindigkeit erforderlich. Diese notwendige Rigidität der Zustandsgleichung beschränkt das mögliche Auftreten von starken Phasenübergängen mit niedrigen Schallgeschwindigkeiten und ausgedehnten Koexistenzregionen in Neutronensternkernen. In Anbetracht dieser Einschränkungen durch die verfügbaren Daten werden viel diskutierte Themen, wie die Frage nach einem Phasenübergang der zur Wiederherstellung der chiralen Symmetrie führt und die aktiven Freiheitsgrade in kalter und dichter baryonischer Materie, untersucht. Zudem entwickeln wir eine neue Methode mit der die Zustandsgleichung direkt aus Detektorbeobachtungen abgeleitet werden kann ohne einen zuvor erforderlichen Zwischenschritt. In dieser Methode wird die analytisch nicht verfügbare Likelihood der Neutronensterndaten durch neuronale simulationsbasierte Inferenzmethoden erlernt. Durch die hieraus resultierende numerische Kostenreduktion für die Einbeziehung zusätzlicher Beobachtungen kann die Inferenz mit dem erwarteten Anstieg an Neutronensterndaten in den kommenden Jahren skalieren.

All the remaining information in Fig. 1 [a sketch showing possible phases of QCD] is based at best on educated conjecture and at worst on wild speculation.

- Larry McLerran [1], 2020

ACKNOWLEDGMENTS

I am profoundly grateful to Prof. Wolfram Weise and Prof. Norbert Kaiser for the opportunity to pursue my PhD under their guidance. Their kind support allowed me to explore fascinating research questions. It was a privilege to collaborate so closely with two such distinguished researchers, and their infectious enthusiasm for physics proved to be a constant source of motivation.

My deepest gratitude goes to Prof. Kenji Fukushima for hosting me during two extended stays at The University of Tokyo. His hospitality, along with the stimulating discussions with him and his group, made the experience truly enriching. With his support, I was able to meet other scientists across Japan.

Special thanks to Prof. Lukas Heinrich, who taught me much about deep learning methods for statistical inference. His introduction to Dr. Aishik Ghosh and Dr. Chirag Modi led to a fruitful collaboration with always changing time zone differences. I thoroughly enjoyed their company and learned a lot from their expertise. My gratitude extends to Prof. Daniel Whiteson and Dr. Delaney Farrell for their valuable collaboration.

Thanks to my colleagues Maurus Geiger and Niklas Döpfer for the many lunchtime discussions in the cafeteria. I wish them much success with their dissertations.

Finally, my deepest appreciation goes to my friends and family for their support throughout this thesis journey. While it is beyond the scope of this acknowledgment to name them all, my special thanks go to Dominik Stamen and Lucas Pache, who also assisted with proofreading this thesis. I am immensely grateful to Anna Fuhrmann for her constant support, especially during the writing of this thesis. By now, she has probably been the most frequent audience of my research presentations.

PUBLICATIONS

This thesis contains material which has previously been published in:

- Len Brandes, Chirag Modi, Aishik Ghosh, Delaney Farrell, Lee Lindblom, Lukas Heinrich, Andrew W. Steiner, Fridolin Weber and Daniel Whiteson, ‘Neural simulation-based inference of the neutron star equation of state directly from telescope spectra’, *JCAP* **09**, 009 (2024).
- Len Brandes and Wolfram Weise, ‘Constraints on phase transitions in neutron star matter’, *Symmetry* **16**, 111 (2024).
- Len Brandes, Wolfram Weise and Norbert Kaiser, ‘Evidence against a strong first-order phase transition in neutron star cores: impact of new data’, *Phys. Rev. D* **108**, 094014 (2023).
- Len Brandes, Wolfram Weise and Norbert Kaiser, ‘Inference of the sound speed and related properties of neutron stars’, *Phys. Rev. D* **107**, 014011 (2023).

CONTENTS

1. INTRODUCTION	1
2. BACKGROUND	5
2.1. Neutron stars	5
2.1.1. Formation & properties	5
2.1.2. Tolman–Oppenheimer–Volkoff equations	8
2.1.3. Tidal deformability	9
2.2. Quantum Chromodynamics	11
2.2.1. Fundamental symmetries	11
2.2.2. Chiral effective field theory	13
2.2.3. Perturbative QCD	14
2.2.4. Phases of strongly interacting matter	17
2.3. Deep learning	19
2.3.1. Artificial neural networks	20
2.3.2. Training	20
2.3.3. Normalizing flows	21
3. INFERENCE OF THE EQUATION OF STATE	25
3.1. Equation of state parametrizations	25
3.1.1. Gaussian	26
3.1.2. Segments	26
3.1.3. Spectral	27
3.2. Bayesian inference	28
3.2.1. Bayes’ theorem	28
3.2.2. Kernel density estimation	30
3.2.3. Credible bands	30
3.2.4. Bayes factors	31
3.3. Priors	31
3.3.1. Parameter priors	32
3.3.2. Mass prior	35
3.3.3. Spectral parametrization	35
3.3.4. Monotonically rising speed of sound	36
3.4. Data & likelihoods	36
3.4.1. Mass measurements	37
3.4.2. Mass-radius measurements	39
3.4.3. Tidal deformability measurements	43
3.4.4. Nuclear data	45
3.4.5. Theory constraints	46

3.5. Simulation-based inference	49
3.5.1. Simulated neutron stars	50
3.5.2. Previous work	52
3.5.3. Bayesian inference with neural likelihood estimation	53
4. RESULTS AND DISCUSSION	57
4.1. Conventional Bayesian inference	57
4.1.1. Parametrization dependence	57
4.1.2. Impact of PSR J0952-0607	64
4.1.3. Evidence for (or against) a strong first-order phase transition	69
4.1.4. Trace anomaly measure	75
4.1.5. Impact of theory constraints	78
4.1.6. Further measurements	82
4.2. Phenomenology and models	85
4.2.1. Low-energy nucleon structure	85
4.2.2. Quark-hadron continuity	88
4.2.3. Chiral symmetry restoration	88
4.2.4. Fermi liquid picture	89
4.3. Inference directly from telescope spectra	92
4.3.1. Example posterior distribution	92
4.3.2. Increasing the number of observations	95
4.3.3. Average performance on test set	96
4.3.4. Discussion & outlook	98
5. SUMMARY AND CONCLUSIONS	101
A. EFFECTIVE ENTHALPY	103
B. ROTATION CORRECTION	105
C. EOS TABLES	107
D. BAYES FACTORS	109
E. HAMILTONIAN MONTE CARLO	111

INTRODUCTION

Neutron stars are among the most extreme objects in the universe. The densities inside their cores far exceed those reachable in terrestrial experiments, so that they provide a unique window into the physics of strongly interacting matter under extreme density conditions [2–6].

Just before the neutron was discovered in 1934, Lev Landau introduced the notion of compact, dense stars that resemble giant atomic nuclei [7]. Only two years later, Walter Baade and Fritz Zwicky advanced the idea that neutron-rich stars could form from the remnants of supernovae following core collapse, coining the term ‘neutron stars’ [8]. In 1939 Richard Tolman [9] and, separately, Robert Oppenheimer and George Volkoff [10] derived the system of differential equations describing the internal structure of neutron stars from the general theory of relativity, now known as the Tolman–Oppenheimer–Volkoff equations. Three decades later, the first observational evidence of a neutron star was discovered by Jocelyn Bell and Antony Hewish [11].

For a long time the composition of neutron star interiors has been subject to speculation. Under the extreme density conditions realized in their cores, matter could exist in various exotic states [12], such as baryons in the form of hyperons and Δ isobars [13–17], deconfined quarks [18–21], color superconducting phases [22–24], quarkyonic matter [25, 26], or possibly meson condensates [27–31]. The situation changed drastically when, via pulsar timing of general relativistic Shapiro time delays, several heavy neutron stars with masses around two times the mass of the sun were discovered [32–37]. These observations set severe constraints on the neutron star matter equation of state (EoS), i.e., the relationship between pressure and energy density, since the pressure inside the stars must be sufficiently high to stabilize such heavy objects against gravitational collapse. Some previously discussed simple forms of exotic matter could, therefore, be excluded when their corresponding EoSs turned out to be too soft. Further important information came from the gravitational wave signals of binary neutron star mergers observed by the LIGO and Virgo Collaborations [38, 39], with resulting constraints on their tidal deformabilities. Moreover, the NICER telescope on board the International Space Station measured pulse profiles of hot spots on the surfaces of rapidly rotating pulsars. From those, the combined mass and radius of the corresponding two neutron stars could be inferred [40–43]. This database has recently been extended by the black widow pulsar PSR J0952-0607 [44], which is the heaviest neutron star found so far and also one of the fastest rotating pulsars. A second recently reported object is the supernova remnant HESS J1731-347 [45] with an unusually small mass and radius. The most recent addition to the database is PSR J0437-4715, the third neutron star measured by the NICER telescope, for which preliminary results have only very recently become available [46]. Many more multimessenger observations are expected in the future.

In principle the physics of dense matter in the core of neutron stars is governed by Quantum Chromodynamics (QCD), the theory of the strong interaction. At asymptotically high baryon densities far beyond those encountered in neutron stars, weakly interacting quarks and gluons are the relevant active degrees of freedom, and perturbative QCD (pQCD) computations become feasible [47–49]. The opposite, low-density limit is accessible by chiral effective field theory (ChEFT) as the low-energy realization of QCD. ChEFT provides a systematic framework with controllable uncertainties for dealing with nuclear many-body systems [50, 51]. This approach gives satisfactory descriptions of nuclear and neutron matter [52] in a range of validity up to densities $n \sim 1 - 2 n_0$, with $n_0 \simeq 0.16 \text{ fm}^{-3}$ the equilibrium density of nuclear matter. Between these extremes, the nature and location of the transition from cold dense nuclear or neutron matter to quark and gluon degrees of freedom is still largely unknown. With the ab-initio method of numerically solving QCD on a lattice at non-zero baryon densities severely hindered by the sign problem of the fermionic determinant [53], theoretical calculations extending from $n \sim 2 n_0$ to high densities are mostly based on models [54]. Various hypotheses have been discussed in the literature, ranging from a first-order (chiral) phase transition to a continuous hadron-quark crossover [25, 26, 47, 50, 55–58]. All possible scenarios for phase transitions or crossovers in dense matter would leave their characteristic signatures in the speed of sound, $c_s^2 = \partial P(\varepsilon)/\partial \varepsilon$. This motivates an investigation of constraints on the behavior of the sound speed in neutron stars based on the presently available data [59].

In this work we translate the recently collected exterior neutron star observables, together with nuclear physics constraints at low densities from ChEFT [60, 61] and constraints from pQCD at asymptotically high densities [49, 62], into restrictions for interior neutron star properties using Bayesian inference, an approach which has been used extensively in the literature [2, 63–107]. There has been a growing interest in complementary methods based on machine learning [108–128]. Alternative approaches simply remove equations of state that do not reproduce neutron star properties within the credible intervals of the astrophysical observables [129–137]. In our analysis we pay particular attention to the speed of sound inside neutron stars. It is modeled using multiple general representations that can describe a variety of different phase structures depending on their parameters [67, 130, 131, 138]. A comparison of inference results with different functional forms as input gives an impression of the model dependence caused by possible biases in the choices of parametrizations. We investigate the roles of low- and high-density constraints from ChEFT and pQCD in detail. From the inferred behavior of the sound speed and the equation of state we can deduce implications regarding the likelihood of phase transitions inside neutron stars. Special emphasis is put on the additional impact of the heaviest known pulsar, PSR J0952-0607, in comparison with previously available data. Part of this discussion also concerns the possible range of validity for a description of neutron star matter in terms of conventional baryonic degrees of freedom. In view of the constraints emerging from the analysis of the observational data, we discuss implications for the structure and composition of dense matter in neutron star cores. We address the quest for a chiral phase transition, quark-hadron continuity and its realization as a soft crossover and a Fermi liquid approach to neutron star matter.

Inferring the equation of state of neutron star matter directly from astrophysical multi-messenger detector data is not numerically feasible using conventional Bayesian methods. Therefore, in our inference analysis we follow the traditional approach to carry out the inference in two steps [139, 140]. In a first step, exterior neutron star properties, such as the mass, radius and tidal deformability, are inferred from the observed detector data. Then, the EoS is inferred from a set of exterior neutron star properties in a second step,

introducing additional assumptions. Recent works have shown promise in directly inferring the EoS from detector data in a single step [127, 128]. These studies analyzed spectroscopic measurements of thermal surface emissions from quiescent low-mass X-ray binaries (qLMXBs) [141] using deep learning methods, where artificial neural networks are trained to model complex non-linear relationships. Similar techniques have successfully been applied to study strongly interacting matter under extreme conditions [142, 143]. While these approaches offer a direct EoS prediction from observed telescope spectra, their high computational cost limits access to the full posterior distribution. This work presents a novel method that addresses the limitations of previous approaches. We employ neural likelihood estimation (NLE) [144], a recently developed simulation-based inference technique that utilizes normalizing flows, a class of neural networks capable of modeling high-dimensional probability distributions [145, 146]. Once trained, the evaluation of normalizing flows 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 sampling methods, such as Hamiltonian Monte Carlo (HMC) [147, 148]. This approach naturally scales to the anticipated increase in neutron star observations in the future, as it does not require retraining to apply to larger data sets.

This thesis is organized as follows: chapter 2 reviews the theoretical background relevant for inferring the neutron star EoS. We discuss the formation and properties of neutron stars, review the theory of gravity necessary for deriving the Tolman–Oppenheimer–Volkoff equations, and provide an introduction to Quantum Chromodynamics and the possible phases of strongly interacting matter at finite temperatures and chemical potentials. Finally, we introduce neural networks and their training, focusing specifically on normalizing flows. In chapter 3, we introduce the parametrizations used to model the speed of sound in neutron stars and explain the statistical procedures employed to infer constraints on neutron star properties. This chapter details both the conventional two-step method as well as our novel simulation-based inference approach. We also provide an overview of the currently available multimessenger observations of neutron stars. Chapter 4 presents the results of our inference. We examine the evidence for (or against) a strong first-order phase transition within neutron stars and discuss implications for the phenomenology of neutron star matter. Lastly, we show the performance of our novel one-step inference method on a test set of simulated data. The final chapter 5 provides a summary of the key findings and concludes this thesis.

Throughout this work, we will be working in natural units, in which both the reduced Planck constant \hbar and the speed of light c are set to unity.

BACKGROUND

In this chapter, we introduce the theoretical background necessary for the inference of the neutron star matter equation of state. We begin by discussing the formation of neutron stars in core-collapse supernovae and explain typical ranges for their macroscopic properties. We briefly introduce the theory of gravity necessary to derive the Tolman–Oppenheimer–Volkoff equations. These equations can be solved for a given neutron star equation of state, or alternatively a given speed of sound, to determine the possible values for the masses and radii of neutron stars. Additional equations can be solved to obtain the tidal deformability, which quantifies the response of a star to an external gravitational field. In order to understand the physics of neutron star matter which determines the equation of state, we review the theory of the strong interaction, Quantum Chromodynamics (QCD). We discuss the fundamental symmetries of strongly interacting matter and introduce computational frameworks available at different density regimes; chiral effective field theory at low densities and perturbative QCD at asymptotically high densities. Furthermore, we provide an overview of the possible phases of strongly interacting matter at finite temperatures and chemical potentials. At the end of this chapter, we introduce fundamental concepts of deep learning, focusing on how artificial neural networks are trained to approximate non-parametric functions. We specifically discuss normalizing flows, a class of neural networks designed to model complex probability distributions.

2.1 NEUTRON STARS

2.1.1 FORMATION & PROPERTIES

Neutron stars are one of the last stages in the stellar evolution of massive stars with a mass exceeding eight solar masses, $M > 8 M_{\odot}$ [149]. The sequential stages culminating in the formation of a neutron star are sketched in Fig. 2.1. Stars remain stable through thermonuclear fusion within their central cores, where thermal and radiation pressure counterbalance gravitational forces. The fusion of hydrogen into helium and subsequently heavier elements continues until the creation of iron, marking the endpoint of exothermic nuclear fusion reactions. Consequently, stellar nucleosynthesis results in the formation of an iron-rich core, which is supported by electron degeneracy pressure. The surrounding shell continues fusion, adding further mass to the core.

When the core mass surpasses the Chandrasekhar limit, $M \gtrsim 1.4 M_{\odot}$, gravitational forces overcome electron degeneracy pressure, prompting core collapse [151]. The collapse leads to a rapid increase in temperature, enabling the creation of neutrons through electron capture, $p + e^{-} \rightarrow n + \nu_e$, which releases a large amount of neutrinos. At some point, the

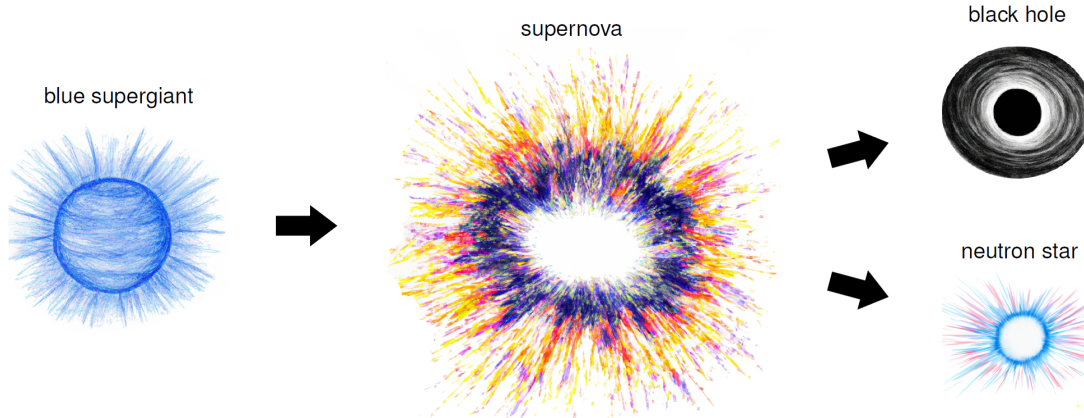


FIG. 2.1. Stages in the stellar evolution of massive stars: through stellar nucleosynthesis an iron-rich core forms in a supergiant star. When the core mass exceeds 1.4 solar masses, it collapses. The implosion is halted by strong force repulsion and neutron degeneracy pressure, and bounces back into a supernova. After most of the material has been ejected, the remaining collapsed core forms either a black hole or a neutron star, depending on its mass [images created with DALL·E, openAI [150]].

contraction is impeded by the repulsion of strong nuclear forces and neutron degeneracy pressure, leading to a rebound and outward bounce. The resultant shock wave combined with the released neutrinos triggers a supernova explosion that expels most of the star's material, leaving behind only the collapsed core. Depending on its mass, the remnant forms either a neutron star or a black hole. Neutron stars can also originate from white dwarfs that collapsed following mass accretion from a companion or through a merger.

Because the angular momentum of the initial star is conserved, but its radius is dramatically reduced in the collapse, the remaining neutron star has a greatly increased rotational speed. The fastest observed neutron stars have orbital periods of only 1.4 ms [44, 152] and correspondingly large angular velocities ω of 4.5×10^3 rotations per second. The velocity at the surface of the neutron star must be less than or equal to the speed of light c ,

$$v_{\text{surf}} = \omega R \leq c . \quad (2.1)$$

This shows that with the increase of ω due to the collapse, the radius R reduces from the star's initial value of about 10^6 km to only $R \sim 10$ km. Conservation of the magnetic flux, $\phi = B\pi R^2$, implies that the magnetic field strength B must increase considerably with the drastic reduction of the radius. Consequently, neutron stars exhibit some of the most intense magnetic fields observed in the universe, reaching values up to 10^{15} G at their surfaces [153]. These strong magnetic fields accelerate particles, leading to the emission of synchrotron radiation. If the star's magnetic- and rotation axis are not aligned, the double cone of emitted radiation rotates with the frequency of the neutron star. If earth lies within this radiation path, we can observe regularly recurring signals from this *pulsar*. This is how neutron stars were first discovered in 1967 [11]. Since then several thousand pulsar sources have been detected [154, 155]. Over time, pulsars gradually spin down due to the conversion of their rotational energy into emitted electromagnetic radiation [4].

For the stars to maintain stability, the gravitational force acting on a test mass m at their surface must be equal or greater than the outwards directed centrifugal force,

$$G_N \frac{Mm}{R^2} = m\omega^2 R , \quad (2.2)$$

where G_N is the gravitational constant. This stability condition indicates that the masses of neutron stars are of the order of the sun’s mass, $M \sim M_\odot$. In fact, stellar evolution calculations of core-collapse supernovae predict a minimum neutron star mass of $1.17 M_\odot$ [156], while the maximum mass for neutron stars is expected to be in the range $M_{\max} \sim 2 - 2.5 M_\odot$. Their extreme masses and small radii result in neutron stars having densities up to six times the equilibrium density of nuclear matter, $n_0 = 0.16 \text{ fm}^{-3}$. These densities far surpass those reachable in terrestrial experiments [157], such that neutron stars provide a unique window to study physics at extremely high densities. At birth in a supernova, neutron stars have temperatures of $10^{11} - 10^{12} \text{ K}$ [149, 158], rapidly cooling to less than 10^6 K through neutrino emission [159]. This temperature is much smaller than the large Fermi temperature of nucleons in neutron stars, which can reach values up to $\sim 10^{12} \text{ K}$, so that we can consider neutron stars to be cold.

The structure of neutron stars is depicted in Fig. 2.2. They are believed to possess a thin gaseous atmosphere extending up to several centimeters [160] which is primarily composed of ionized hydrogen, with the potential inclusion of helium or carbon. While its mass is negligible, the atmosphere influences thermal emission emitted from the neutron star’s surface [161]. Beneath the surface lies the neutron star crust, reaching a depth of around 1 kilometer [160]. The outer crust is comprised of ionized atoms, forming a lattice of iron nuclei surrounded by a degenerate relativistic gas of electrons. As the density increases, the electron Fermi energy rises, making electron capture energetically favorable and leading to the formation of increasingly neutron-rich nuclei. Around a neutron drip density of approximately $2 \times 10^{-3} n_0$, neutrons start dripping out from the nuclei, forming clusters of nuclear matter in the inner crust [3]. With a further increase of density, these clusters cannot retain their spherical form due to the interplay between nuclear and electromagnetic forces, such that they take complex shapes commonly referred to as ‘nuclear pasta’ [161]. The neutron star crust extends to a density of around $0.5 n_0$, before transitioning to a state of pure nuclear matter in the neutron star core [160]. This core, constituting the majority of the neutron star’s mass, is predominantly composed of neutrons, which are in beta equilibrium with a small fraction of protons and electrons to ensure charge neutrality. Muons begin to appear when their chemical potential equals that of the electrons. The innermost core of neutron stars may potentially exhibit exotic degrees of freedom as detailed in Sec. 2.2.4.

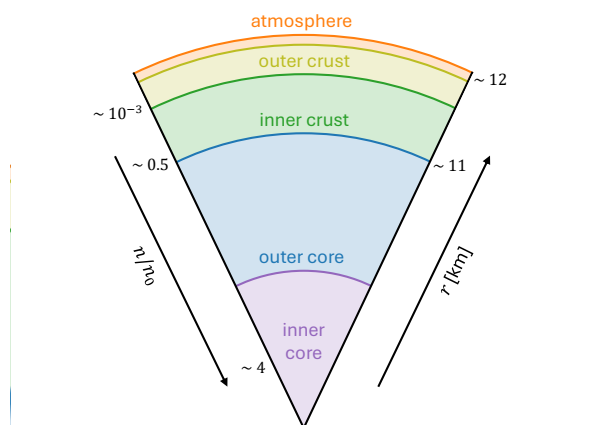


FIG. 2.2. Schematic cross-section of a neutron star. With decreasing radius r the baryon density n inside the neutron star (in units of the nuclear saturation density, $n_0 = 0.16 \text{ fm}^{-3}$) increases. The radii are not true to scale.

2.1.2 TOLMAN–OPPENHEIMER–VOLKOFF EQUATIONS

The theory of general relativity, developed by Albert Einstein [162], provides a geometric framework for understanding gravitation as the curvature of four-dimensional spacetime induced by mass and energy. Objects move on curved paths influenced by the geometry of spacetime, leading to phenomena such as the bending of light by massive objects apparent from gravitational lensing. This theory is needed to understand the internal structure of neutron stars, where extreme conditions of density and gravity necessitate a departure from classical Newtonian gravity. In this section we briefly introduce the theory of gravity necessary for the derivation of the Tolman–Oppenheimer–Volkoff equations¹, for more details see, e.g., Refs. [164, 165].

Let us denote a general spacetime coordinate as x^μ . The curvature of spacetime can be represented via a metric $g_{\mu\nu}$, allowing the definition of an invariant distance as

$$ds^2 = g_{\mu\nu} dx^\mu dx^\nu . \quad (2.3)$$

The fundamental field equation of general relativity is Einstein’s equation:

$$G_{\mu\nu} = 8\pi G_N T_{\mu\nu} , \quad (2.4)$$

where the Einstein tensor $G_{\mu\nu}$, derived from the metric tensor $g_{\mu\nu}$ and its derivatives, characterizes the spacetime curvature. In flat spaces without gravity, i.e., in the absence of masses, the Einstein tensor becomes zero. For finite masses, it is proportional to the symmetric energy-momentum tensor $T_{\mu\nu}$, which adheres to energy-momentum conservation,

$$\nabla_\mu T^{\mu\nu} = 0 . \quad (2.5)$$

In the presence of a spherically symmetric and static source, such as a neutron star, the metric takes the following form:

$$ds^2 = \alpha(r) dt^2 - \beta(r) dr^2 - r^2 d\Omega , \quad (2.6)$$

where the radius-dependent functions $\alpha(r)$ and $\beta(r)$ can be determined by taking the Newtonian limit of small masses and weak gravity. Outside the source, the energy momentum tensor vanishes and the metric takes the form of the Schwarzschild solution. Assuming matter inside the neutron star can be described as an ideal fluid with isotropic pressure P , four-velocity u_μ and energy density ε , the energy momentum tensor takes the following form:

$$T_{\mu\nu} = (\varepsilon + P) u_\mu u_\nu - P \cdot g_{\mu\nu} . \quad (2.7)$$

By matching to the exterior Schwarzschild solution, one can derive the coupled system of differential equations called the Tolman–Oppenheimer–Volkoff (TOV) equations [9, 10]

$$\frac{\partial P(r)}{\partial r} = -\frac{G_N}{r^2} [\varepsilon(r) + P(r)] \left[m(r) + 4\pi r^3 P(r) \right] \left(1 - \frac{2G_N m(r)}{r} \right)^{-1} , \quad (2.8)$$

$$\frac{\partial m(r)}{\partial r} = 4\pi r^2 \varepsilon(r) . \quad (2.9)$$

Given a relationship between pressure and energy density, $P(\varepsilon)$, commonly referred to as the *equation of state*, this coupled system of differential equations can be solved with the boundary condition $m(r=0) = 0$ and a given central pressure $P(r=0) = P_c$.

¹ Parts of the text in this section have been adapted from Refs. [96, 163].

The total mass of the star can be computed by integrating the second TOV equation $M = m(R) = 4\pi \int_0^R dr r^2 \varepsilon(r)$, while the star's radius R is determined as the point at which the pressure vanishes, i.e., $P(R) = 0$. Directly solving to find the pressure vanishing point is not convenient for numerical solutions. Instead, following [166], an approach based on the effective enthalpy $h = \int_0^P dP' (\varepsilon(P') + P')^{-1}$ can be employed, see Appendix A for more details. The TOV equations describe non-rotating neutron stars, with the impact of rotation on M and R expected to become only relevant for very high pulsar spin frequencies [161, 167]. Numerical solutions of the TOV equations for various central pressures yield possible combinations of M and R , forming the *mass-radius relation*. Stability of a sequence of masses and radii is guaranteed only when the neutron star masses increase with the central pressures,

$$\frac{dM}{dP_c} > 0, \quad (2.10)$$

generally culminating in a *maximum supported mass* M_{\max} , beyond which no stable solutions exist. This endpoint of the mass-radius relation is linked to a maximum central pressure, $P_{c,\max}$, and a corresponding maximum energy density, $\varepsilon_{c,\max}$, reached in the star's center. In general, stiffer equations of state reaching higher pressures at smaller energy densities tend to support larger maximum masses. Extended first-order phase transitions can lead to disconnected mass-radius relations, where the endpoint of the (M, R) -sequence does not necessarily correspond to the maximum supported mass.

Matter in the interior of a neutron star can be described in terms of the squared adiabatic speed of sound,

$$c_s^2(\varepsilon) = \frac{\partial P(\varepsilon)}{\partial \varepsilon} \geq 0, \quad (2.11)$$

from which the equation of state can be determined via

$$P(\varepsilon) = \int_0^\varepsilon d\varepsilon' c_s^2(\varepsilon'). \quad (2.12)$$

Causality demands that the speed of sound must always remain smaller than or equal to the speed of light, i.e., $c_s \leq 1$ in units where c is set to unity. In addition, thermodynamic stability of the star dictates that the derivative $\partial P/\partial \varepsilon$ must be non-negative. As a measure of the coupling strength of matter, the behavior of the sound speed c_s^2 is a prime indicator for the phase structure realized inside neutron stars, as will be explained in more detail in Sec. 2.2.4.

2.1.3 TIDAL DEFORMABILITY

In addition to its mass and radius, we can also gain insights into the internal structure of a neutron star by probing its response to the gravitational field of a nearby massive body, e.g., a second neutron star. In such binary systems, the external quadrupolar tidal field of the companion star ε_{ij} induces a quadrupole moment in the neutron star Q_{ij} [168, 169],

$$Q_{ij} = -\lambda \varepsilon_{ij}. \quad (2.13)$$

Here, λ is the tidal deformability, which is related to the $l = 2$ tidal Love number $k_2 = 3/2 \lambda R^{-5}$ that quantifies the susceptibility of a celestial body to deform under the influence of external tidal forces (to leading order in spherical harmonics). The external

tidal field ε_{ij} and the star's quadrupole moment Q_{ij} are defined as coefficients in an asymptotic multipole expansion of the metric at large distances. In the Newtonian limit ε_{ij} is equivalent to the second spatial derivative of the external gravitational potential [160, 170].

The change from equilibrium due to an external quadrupolar field can be modeled via a linear static metric perturbation $g_{\mu\nu} = g_{\mu\nu}^{(0)} + h_{\mu\nu}$ in the Einstein equation (2.4). The components of the perturbation $h_{\mu\nu}$ can be expressed in terms of a function $H(r)$ [169]. This results in the following system of coupled differential equations that have to be solved simultaneously with the TOV equations (2.8 - 2.9):

$$\frac{\partial H}{\partial r} = \beta, \quad (2.14)$$

$$\begin{aligned} \frac{\partial \beta}{\partial r} = & \frac{2\beta}{r} \left(1 - \frac{2m}{r}\right)^{-1} \left\{ -1 + \frac{m}{r} + 2\pi r^2(\varepsilon - P) \right\} + 2 \left(1 - \frac{2m}{r}\right)^{-1} H \\ & \times \left\{ -2\pi \left[5\varepsilon + 9P + (\varepsilon + P) \frac{\partial \varepsilon}{\partial P} \right] + \frac{3}{r^2} + 2 \left(1 - \frac{2m}{r}\right)^{-1} \left(\frac{m}{r^2} + 4\pi r P \right)^2 \right\}. \end{aligned} \quad (2.15)$$

This system has to be integrated outwards starting just outside the center using the expansion $H = a_0 r^2$ and $\beta = 2a_0 r$ for small radii r [171]. The constant a_0 can be chosen arbitrarily as it cancels out in the later expression for k_2 . After defining the function $y(R) = R \beta(R)/H(R)$ as well as the compactness $C = M/R$, the $l = 2$ tidal Love number can be determined from:

$$\begin{aligned} k_2 = & \frac{8C^5}{5} (1 - 2C^2)(2 + 2C(y - 1) - y) \left\{ 2C [6 - 3y + 3C(5y - 8)] \right. \\ & \left. + 4C^3 [13 - 11y + C(3y - 2) + 2C^2(1 + y)] + 3(1 - 2C)^2 \right. \\ & \left. \times [2 - y + 2C(y - 1)] \log(1 - 2C) \right\}^{-1}. \end{aligned} \quad (2.16)$$

From k_2 we can compute the *dimensionless tidal deformability*² [171]

$$\Lambda = \frac{2}{3} k_2 \left(\frac{R}{M} \right)^5. \quad (2.17)$$

With the outlined procedure, for a given neutron star equation of state the corresponding tidal deformability can be computed together with the mass-radius relation. This is illustrated in Fig. 2.3 for one example EoS. Here, we use the `LALSuite` library for a fast numerical solution of the full system of coupled differential equations [172].




FIG. 2.3. Numerically solving the TOV equations (2.8 - 2.9) together with Eqs. (2.14) and (2.15) for various central pressures yields the mass-radius relation (M, R) as well as the tidal deformability Λ for a given equation of state, $P(\varepsilon)$, illustrated here for one example EoS.

² In the following, we will just use the term ‘tidal deformability’ to refer to the dimensionless tidal deformability Λ .

2.2 QUANTUM CHROMODYNAMICS

In this section we give a brief introduction to the theory of the strong interaction, which has been formulated on the basis of the fundamental symmetries that govern the dynamics between quarks and gluons³. One important aspect of the strong interaction is its approximate chiral symmetry, which serves as a guiding principle for a systematic low-energy description via chiral effective field theory. In contrast, the behavior at high energies can be understood through perturbative quantum field theory calculations. Strongly interacting matter is characterized by a complex phase structure at finite temperatures and chemical potentials. The primary degrees of freedom at high densities are still unknown, but insights can be gained from the study of neutron stars.

2.2.1 FUNDAMENTAL SYMMETRIES

The standard model of particle physics stands as the preeminent theoretical framework describing the fundamental constituents of matter and their interactions, excluding the gravitational force (which is described by the theory of general relativity, see Sec. 2.1.2). A central part of this model is Quantum Chromodynamics (QCD), the quantum field theory of the strong interaction that describes the dynamics between quarks and gluons [173].

Quarks are fermionic particles with a spin of 1/2 whereas gluons are massless bosons with a spin of 1. There are six different flavors of quarks — up, down, charm, strange, top, and bottom — each endowed with distinct masses [174]. The strong interactions between quarks and gluons are responsible for the formation and processes of hadrons, composite particles made of two or more quarks. The discovery of hadrons like Δ^{++} and Ω^- baryons that appear to consist of three quarks with the same quantum numbers was historically important for the development of QCD. This apparent violation of the Pauli exclusion principle necessitated the introduction of a novel quantum number: color.

Color, in the context of QCD, is an attribute similar to the electric charge, albeit with three different manifestations: red, green, and blue, as well as their anti-colors. With this additional quantum number to distinguish quarks, the apparent paradox of same-flavor baryons could be resolved. Unlike the electromagnetic force, wherein charge neutral photons mediate interactions, gluons, as carriers of the strong force, carry color charges. This characteristic manifests in self-interactions among gluons. Both gluon self-interactions and the interactions between quarks and gluons exhibit invariance under $SU(3)_c$ rotations in color space. Guided by this fundamental gauge symmetry the Lagrangian of Quantum Chromodynamics (QCD) has been constructed, providing a mathematical description of the underlying dynamics of strong interactions:

$$\mathcal{L}_{\text{QCD}} = \sum_f \bar{q}_f (iD_\mu \gamma^\mu - m_f) q_f - \frac{1}{4} F_{\mu\nu}^a F^{a\mu\nu} . \quad (2.18)$$

Here, q_f represents the quark fields, while m_f signifies the corresponding masses for each quark flavor f , and γ_μ are the Dirac gamma matrices. The covariant derivative D_μ is given by

$$D_\mu = \partial_\mu + ig A_\mu^a(x) T^a . \quad (2.19)$$

In addition to the kinetic term of the quarks, this characterizes the interaction of quarks with gluons, with A_μ^a representing the gluonic fields, g being the strong coupling constant,

³ Parts of the text in this section have been adapted from Refs. [96, 163].

and T^a denoting the generators of the $SU(3)_c$ group. The latter can be explicitly represented in terms of the Gell-Mann matrices, $T^a = \lambda^a/2$. The gluon field-strength tensors $F_{\mu\nu}^a$, are defined as:

$$F_{\mu\nu}^a = \partial_\mu A_\nu^a(x) - \partial_\nu A_\mu^a(x) - gA_\mu^b(x)A_\nu^c(x)f^{abc} . \quad (2.20)$$

Here, f^{abc} are the totally antisymmetric structure constants associated with the Lie algebra formed by the generators, $[T^a, T^b] = if^{abc}T^c$.

Due to the three- and a four-gluon self-interactions, which arise from the product of the two field-strength tensors in the Yang–Mills term in Eq. (2.18), we do not observe any isolated color-charged particles. Quarks and gluons are always contained in color-neutral composite objects such as hadrons. This phenomenon, referred to as confinement, is manifest in the coupling constant of the strong interaction, which becomes increasingly large at small energies. Consequently, perturbative expansions in the coupling constant, which have been very successful in computing electroweak interactions, are not feasible for quark and gluon interactions in this energy regime.

Instead, in a strategy known as lattice QCD, observables can be numerically computed using Monte Carlo methods [175]. At large baryon densities and small temperatures this approach is severely hindered by the sign problem [53]; for non-zero chemical potentials the fermionic determinant acquires a complex phase, so that it can no longer be interpreted as a probability weight in the Monte Carlo procedure. Alternatively, non-perturbative techniques, including Dyson–Schwinger equations and functional renormalization group methods, have been developed [176, 177]. Beyond these approaches, low-energy descriptions capturing the relevant degrees of freedom, i.e., baryons and mesons, can be employed to describe strongly interacting matter in this energy regime, see Sec. 2.2.2. The strong coupling decreases at higher energies, making perturbative calculations of QCD feasible under conditions of extremely high temperatures or densities, as explained in Sec. 2.2.3. This behavior is termed asymptotic freedom.

An important characteristic of the QCD Lagrangian lies in its approximate *chiral symmetry*. This becomes evident through the projection of quark fields onto their left- and right-handed components, denoted as $q_{L/R} = P_{L/R}q$, with projectors $P_{L/R} = \frac{1}{2}(1 \mp \gamma_5)$ employing the Dirac gamma matrix γ_5 . Under these projections, the QCD Lagrangian takes the following form:

$$\mathcal{L}_{\text{QCD}} = \bar{q}_L iD^\mu \gamma_\mu q_L + \bar{q}_R iD^\mu \gamma_\mu q_R - \bar{q}_L m q_R - \bar{q}_R m q_L - \frac{1}{4} F_{\mu\nu}^a F^{a\mu\nu} . \quad (2.21)$$

This reformulation reveals that only the mass terms couple left- and right-handed fields. This is of some consequence since the masses of the three lightest quarks (up, down and strange) are small enough to be neglected under most circumstances [174]. In the chiral limit of vanishing light quark masses, the QCD Lagrangian of the light flavors becomes invariant under separate unitary transformations to the left- and right-handed quark fields. This symmetry can be further decomposed into:

$$U(3)_L \times U(3)_R \cong SU(3)_L \times SU(3)_R \times U(1)_L \times U(1)_R . \quad (2.22)$$

According to Noether’s theorem, continuous symmetries yield conserved currents. For the above symmetries, we can combine the conserved left- and right-handed currents into vector ($V = L + R$) and axial ($A = L - R$) currents,

$$V_\mu^a = \bar{q} \gamma_\mu T^a q , \quad A_\mu^a = \bar{q} \gamma_\mu \gamma_5 T^a q , \quad (2.23)$$

$$V_\mu = \bar{q} \gamma_\mu q , \quad A_\mu = \bar{q} \gamma_\mu \gamma_5 q . \quad (2.24)$$

The singlet axial current A_μ is not conserved due to the anomalous breaking of the $U(1)_A$ symmetry in QCD, which is evident from the observed decay of neutral pions into two photons. On the other hand, the singlet vector current V_μ always remains conserved and is linked to baryon number conservation.

The absence of parity doublets in the observed particle spectrum indicates that the axial symmetry, corresponding to the current A_μ^a , is not present in the ground state of the theory. This indicates that the chiral symmetry $SU(3)_L \times SU(3)_R$ is spontaneously broken to $SU(3)_V$. Goldstone's theorem states that there exists one corresponding massless Nambu–Goldstone boson for each generator of a spontaneously broken symmetry. In the context of spontaneously broken chiral symmetry, the eight massless Goldstone bosons can be identified with the lightest mesons; the three pions (π^\pm, π^0), the four kaons (K^\pm, K^0, \bar{K}^0) and the eta meson (η). The non-vanishing masses of the three lightest quarks — m_u, m_d and m_s — explicitly break chiral symmetry, such that the lightest mesons possess finite masses instead of being massless. Since the mass of the strange quark is much larger than the up and down quark masses, kaons and eta mesons are heavier than pions [174]. In the limit $m_u = m_d \neq 0$ the QCD Lagrangian remains invariant under the reduced vector symmetry $SU(2)_V$ corresponding to isospin symmetry.

2.2.2 CHIRAL EFFECTIVE FIELD THEORY

Guided by the symmetry properties of QCD discussed in the preceding section, we can convert the highly non-perturbative structure of the theory into a more transparent framework. *Chiral effective field theory* (ChEFT) provides a systematic ansatz to describe low-energy QCD dynamics, particularly the behavior of light mesons, based on the spontaneously broken chiral symmetry [50, 178, 179]. To construct this effective field theory, Steven Weinberg suggested to formulate the most general Lagrangian consistent with the present symmetries [180]. This entails ensuring that the Lagrangian is invariant under chiral symmetry $SU(3)_L \times SU(3)_R$ while its ground state must remain invariant under the vector symmetry $SU(3)_V$.

Note that the Goldstone bosons transform non-linearly under chiral transformations. Therefore, to construct a chirally invariant Lagrangian, we employ instead the unitary matrix $U(x) = \exp(i\phi(x)/f_0)$ as the central building block, where the Goldstone bosons are collected into the matrix $\phi(x)$. This unitary matrix transforms under chiral transformations as $U(x) \rightarrow RU(x)L^\dagger$, where $(L, R) \in SU(3)_L \times SU(3)_R$. The Lagrangian can then be constructed based on combinations of the field $U(x)$ and its derivatives, with Lorentz invariance allowing only even numbers of derivatives. Expanding in the order of derivatives yields the following effective Lagrangian:

$$\mathcal{L}_{\text{eff}} = \delta\mathcal{L}_M + \mathcal{L}_2 + \mathcal{L}_4 + \dots \quad (2.25)$$

Here the constant zeroth-order term \mathcal{L}_0 was omitted. The contribution $\delta\mathcal{L}_M$ accounts for the explicit symmetry breaking and is linear in the quark mass matrix $M = \text{diag}(m_u, m_d, m_s)$, transforming as $M \rightarrow RML^\dagger$ under chiral transformation. At leading order, the Lagrangian takes the following form:

$$\begin{aligned} \mathcal{L}_{\text{eff}}^{\text{LO}} &= \delta\mathcal{L}_M + \mathcal{L}_2 \\ &= \frac{f_0^2}{2} v \text{Tr}[MU^\dagger + UM^\dagger] + \frac{f_0^2}{4} \text{Tr}[\partial_\mu U \partial^\mu U^\dagger] . \end{aligned} \quad (2.26)$$

Through matching to QCD matrix elements, the constant f_0 is identified as the pion decay constant f_π in the chiral limit, while the constant v is related to the chiral quark condensate.

According to the Lagrangian \mathcal{L}_{eff} , couplings between Goldstone bosons inherently involve derivatives. Consequently, both higher-order terms with more derivatives and loop corrections that introduce additional vertices depend on the squared momenta of the Goldstone bosons. At low energies, these momenta are of order of the Goldstone boson masses, which are small compared to the chiral scale, $\Lambda_\chi \sim 1 \text{ GeV}$. Thus, chiral effective field theory provides a systematic perturbative expansion in the squared momenta and masses of the Goldstone bosons. One can estimate the magnitude of the missing contributions from higher orders using these expansion parameters. When going to higher orders, it becomes necessary to match further low-energy constants that parameterize unresolved short-distance dynamics to experimental data or lattice QCD calculations.

The extension of the effective Lagrangian to include baryons can be done similarly guided by the spontaneously broken chiral symmetry [178, 179]. However, due to the sizable baryon masses of the same order of magnitude as the chiral scale, a revision of the naive power counting of the perturbative series is necessary [181, 182]. The low-energy constants of this effective field theory are fitted to a large amount of empirical nucleon-nucleon and pion-nucleon interaction data. The nuclear two- and three-body forces computed within the ChEFT framework can be combined with many-body methods for an extension to finite densities [51, 183, 184]. These many-body calculations yield compelling results for low-density nuclear phenomenology, i.e., reproducing key features such as the location of the critical point of the nuclear liquid-gas phase transition of symmetric nuclear matter in close agreement with empirical values [52, 185].

At its current state of development ChEFT is believed to be applicable up to densities of approximately $n \sim 1 - 2 n_0$. The results obtained for symmetric nuclear and pure neutron matter within this framework can be extrapolated to neutron star matter by including beta equilibrium conditions [129]. In recent works a novel ansatz has been introduced based on a Gaussian process to derive combined uncertainties from many-body approximations and from missing higher order terms beyond next-to-next-to-next-to-leading chiral order (N³LO) [186–188]. In this way posterior credible bands could be derived for the squared speed of sound, $c_s^2(n)$, as a function of density up to two times nuclear saturation density, $n = 2 n_0$ [60, 61], which are depicted in Fig. 2.4. These credible bands also agree with the results found independently by another group [189]. The ChEFT results serve as a rigorous constraint on the equation of state for neutron star matter at low densities [60, 183, 184, 190]. In Sec. 3.4.5 we outline the implementation of this constraint in our inference analysis. Note that these calculations do not include hyperons, i.e., baryons containing one or more strange quarks, or Δ baryons which might become relevant at higher densities.

2.2.3 PERTURBATIVE QCD

At extremely high densities and temperatures, the coupling constant of the strong interaction decreases due to asymptotic freedom, such that a perturbative weak-coupling expansion of QCD becomes feasible. To incorporate finite temperatures and chemical potentials, one can employ a thermal field theory approach [191]. In the static case of thermal equilibrium, the pressure can be computed from the grand-canonical partition function,

$$Z = \text{Tr}[e^{-\beta(H-\mu N)}], \quad (2.27)$$




FIG. 2.4. Squared sound speed in neutron star matter from N³LO ChEFT results displayed as a function of baryon density n in units of the nuclear saturation density, $n_0 = 0.16 \text{ fm}^{-3}$. Correlated uncertainty bands are given at the 1σ (dashed) and 2σ (dotted) levels. Figure adapted from [61].

where H is the QCD Hamiltonian, N is the particle number operator, $\beta = 1/T$ is the inverse temperature, and μ is the chemical potential ensuring particle number conservation. Central to this approach is the Matsubara formalism, which recognizes the resemblance between Z and the standard QCD partition function, if one identifies the inverse temperature β with an imaginary time $-i\tau$, effectively performing a Wick rotation from Minkowski to Euclidean space. The partition function can be transformed into a path integral over bosonic (ϕ) and fermionic (ψ) fields:

$$Z = \int \mathcal{D}[\psi, \psi^\dagger, \phi] \exp\left(-\int_0^\beta d\tau \int d^3x [\mathcal{L}_E - \mu N]\right), \quad (2.28)$$

where \mathcal{L}_E is the QCD Lagrangian of Eq. (2.18) in Euclidean space. This formalism compactifies the imaginary time dimension τ onto a circle, with periodic and anti-periodic boundary conditions for bosons and fermions, respectively. In the limit of zero temperature, the τ dimension becomes infinite again.

The leading-order term of the pressure derived from this partition function corresponds to that of a relativistic quark gas, i.e., $P \propto \mu^4$. However, a naive loop expansion of the pressure is hindered by infrared divergences [192]. In the medium, gluons become screened and acquire a finite mass through self-energy corrections. Consequently, soft contributions in the gluon propagator need to be resummed to compute the gluon self-energy. This resummation, dominated by the high-momentum behavior of loop corrections, is termed *hard thermal loops* (HTL) [193] and yields contributions logarithmic in the coupling constant. In an HTL effective theory, instead of a naive power expansion in the coupling constant, diagrams are categorized based on their momentum regions, distinguishing between hard modes ($p \gtrsim \mu$), soft modes ($p \lesssim g\mu$), and their interactions within the HTL framework, along with mixed terms from interactions between soft and hard modes. The next-to-next-to-leading-order (N²LO) results at zero temperature and finite density were computed in 1977 for massless quarks [194, 195] and were later extended to include effects of the strange quark mass [196] (see also Ref. [197] for a simplified fit to these results). More recently, utilizing the hard thermal loop prescription, computations have been extended to include soft [49, 198] and mixed [199] contributions at N³LO, although substantial efforts are still necessary to compute the four-loop Feynman diagrams of the hard contributions.

The results obtained at partial N³LO are depicted in Fig. 2.5. The pressure still depends on a scale parameter $X = 3\bar{\Lambda}/2\mu$, proportional to the energy scale $\bar{\Lambda}$ where renormalization conditions are assigned to remove ultraviolet divergences. X is typically allowed to vary within the range of $[1/2, 2]$. The applicability of pQCD results is generally limited to chemical potentials above $\mu \sim 2.4 - 2.6$ GeV, corresponding to densities $n \gtrsim 40 n_0$ well beyond those encountered in neutron stars. Below this threshold, the uncertainty stemming from the indeterminate renormalization scale becomes prohibitively large. As pointed out recently [62, 86], even though pQCD is applicable only at extremely high densities, demanding that any valid EoS should be casually connected to the asymptotic pQCD results can lead to constraints at much lower densities. We include the constraint from pQCD in our inference procedure following the prescription of these works, see Sec. 3.4.5 for more details.




FIG. 2.5. Perturbative QCD pressure from N²LO results including N³LO soft contributions displayed as a function of baryon chemical potential μ . The shaded region shows a renormalization scale variation of $X \in [1/2, 2]$, while the solid line corresponds to $X = 1$. Figure adapted from [49, 62, 86].

In particular, the pQCD results imply that at asymptotic densities the squared speed of sound approaches the *conformal bound*,

$$c_s^2 \rightarrow 1/3, \quad (2.29)$$

from below [197]. This limit can be derived from naive dimensional analysis and asymptotic freedom [200]. In fact it is expected that this bound holds in all conformal theories, i.e., field theories in which the trace of the energy-momentum tensor vanishes [138, 201]. However, recent analyses based on astrophysical observables suggest that this conformal bound can be violated inside neutron stars [71, 72, 74, 86, 132, 202]. Squared sound velocities with $c_s^2 > 1/3$ were also found in recent lattice QCD computations for $N_c = 2$ [203] or at finite isospin chemical potentials [204, 205], where no sign problems are encountered. A possible mechanism for the violation of this bound, based on the trace anomaly in strongly coupled matter, is discussed in Ref. [206]. In that context it is interesting that hard dense loop resummation methods (resumming both the quark and the gluon sector) indicate that the conformal limit may be approached asymptotically from above, with $c_s^2 > 1/3$ [207]. In a later section (4.1.5) we shall examine whether this changed asymptotic behavior has an impact on the sound speed at neutron star core densities.

2.2.4 PHASES OF STRONGLY INTERACTING MATTER

Strongly interacting matter exhibits a rich phase structure at finite temperatures and baryon chemical potentials [47, 208]. Despite considerable theoretical and experimental efforts, many of the intricacies of the QCD phase diagram, especially at large chemical potentials, are still unresolved. The current state of knowledge is sketched in Fig. 2.6.

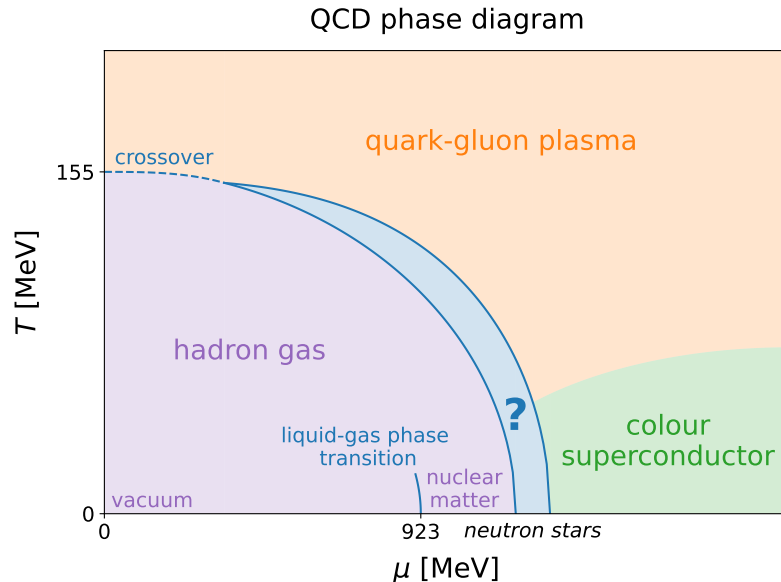


FIG. 2.6. Schematic view of a hypothetical QCD phase diagram. The blue question mark indicates the unknown phase structure at high densities.

At high temperatures and small baryon chemical potentials the phase structure of the strong interaction is well understood from lattice QCD [209, 210] and from high-energy heavy ion collisions [211, 212]. At vanishing baryon chemical potential a continuous crossover proceeds from the hadronic to the quark-gluon phase around a pseudocritical temperature of about 155 MeV. This behavior is reflected in the speed of sound which grows rapidly with increasing temperature in the hadronic phase and then decreases along the crossover transition. At asymptotically high temperatures c_s^2 increases again to reach the asymptotic value of the conformal limit from below (see Fig. 3.3 in Ref. [213]).

At low baryon densities it is quite well established that the thermodynamics of (isospin symmetric) nuclear matter features a first-order liquid-gas phase transition, with a critical point located empirically [185] at a temperature $T_{\text{crit}} \simeq 18$ MeV and density $n_{\text{crit}} \simeq n_0/3$. When viewed in a (T, μ) phase diagram, the first-order liquid-gas transition line starting at the critical point reaches the $T = 0$ axis at a baryon chemical potential $\mu = m_N - B \simeq 923$ MeV corresponding to the binding energy per particle $B \simeq 16$ MeV of symmetric nuclear matter. With an empirical symmetry energy $S \simeq 32$ MeV, this phase transition is absent in pure neutron matter.

At asymptotically high densities and small temperatures, quark and gluon degrees of freedom take over in a *color superconducting phase* [22, 24, 48, 214]. Due to the attractive force between quarks with different colors, quark degrees of freedom near the Fermi surface pair up into diquarks. These diquarks, analogous to Cooper pairs in conventional superconductors, form a condensate which spontaneously breaks $SU(3)_c$ symmetry and leads to finite gluon masses. It is believed that the color-flavor locked (CFL) superconducting

phase represents the energetically favored ground state for systems with three quark flavors at extremely high chemical potentials [23, 47]. At lower densities, where the strange quark mass can no longer be neglected, other less symmetrically paired forms of quark matter may predominate [23, 215].

The detailed nature and density range of the transition from nuclear to quark matter is still unknown, as illustrated by the blue question mark in Fig. 2.6. Many models have been designed, with a variety of hypotheses predicting different active degrees of freedom in this intermediate region. With their low temperatures and core densities well above the nuclear saturation density, neutron stars are the objects of choice to gain information about this speculative region of the phase diagram.

Studies based on Nambu–Jona-Lasinio type models in mean-field approximation have commonly found a first-order chiral phase transition at quite moderate baryon densities for $T = 0$ [25]. A *first-order phase transition* with Maxwell construction is characterized by a region of constant pressure over an interval of density (or energy density) in which two phases coexist [216]. Between the lower and upper endpoints of this phase coexistence interval, the squared sound velocity $c_s^2 = \partial P / \partial \varepsilon$ jumps to zero and back. This is schematically illustrated in Fig. 2.7, which shows typical patterns of phase transitions as they would show up in the sound speed. Sufficiently strong first-order phase transitions can lead to mass-radius relations with a disconnected third-family branch of compact stars containing exotic matter, commonly referred to as twin stars [217]. In a first-order phase transition with Gibbs construction, the pressure in the mixed phase is not constant, but $\partial P / \partial \varepsilon$ still changes discontinuously [218], resulting in a non-zero sound speed. Investigations using non-perturbative functional renormalization group techniques [219–221] found that fluctuations tend to convert the first-order chiral transition present in mean-field studies into a *crossover* shifted to much higher baryon densities, even beyond those realized in neutron star cores (see the discussion in Sec. 4.2.3 for more details). In any case, to support the observed heavy neutron stars with masses $M \sim 2 M_\odot$, a transition to quark matter in neutron stars at relatively low densities is possible only if the quark EoS is extremely stiff, or otherwise the transition has to take place at high densities leading to small quark cores [57, 222].

A possible first-order phase transition at large chemical potentials would need to have a critical endpoint (CEP) to be consistent with the crossover at zero chemical potential and finite temperatures. Extensive efforts are underway in heavy-ion collision experiments to search for signs of such a critical endpoint [157, 223]. Lattice QCD calculations are restricted at finite chemical potentials due to the fermionic sign problem. However, they can be extended to $\mu/T \lesssim 2 - 3$ by employing Taylor expansions of physical quantities around $\mu = 0$ and through analytic continuations of numerical simulations at imaginary chemical potentials. In this constrained domain, no evidence of criticality is observed [224, 225]. Non-perturbative analyses using Dyson–Schwinger equations and functional renormalization group methods agree with lattice results at small chemical potentials but both predict a critical endpoint within a vicinity around $(T_{\text{CEP}}, \mu_{\text{CEP}}) \sim (110, 600)$ MeV [226–229]. Holographic QCD calculations predict similar critical endpoint values [230, 231]. A scenario is also conceivable in which the first-order line starting from the CEP ends in a second critical endpoint at low temperatures and high chemical potentials.

Models proposing a continuous crossover from hadronic to quark matter are often referred to under the key word *quark-hadron continuity*. Such models describe the low-density part of the EoS in agreement with ChEFT calculations but still provide the necessary stiffness to support heavy neutron stars, usually by introducing strongly repulsive correlations in the quark sector [3, 232–235]. A continuous crossover might be realized through an




FIG. 2.7. Characteristic behaviors of the squared sound speed in the presence of a first-order phase transition or a crossover.

intermediate phase of quarkyonic matter [1, 26, 55], a combined phase of quarks and nucleons derived from large N_c considerations [47, 236]. This phase is characterized by a distinct Fermi sea containing both baryons and quarks [26, 99]. At high densities, the Fermi momentum of the baryons increases substantially. Consequently, intermediate momentum states are Pauli blocked, such that particle-hole excitations originating from deep inside the Fermi distribution require substantial energy and momentum transfer. As the strong coupling decreases at large momenta, these degrees of freedom are weakly interacting, pointing to the presence of asymptotically free quarks [99]. Conversely, excitations near the Fermi surface require only small momentum transfers resulting in strong confining forces, so that we might identify these degrees of freedom with baryons. The continuous transition from baryonic to quarks degrees of freedom may be visible as a maximum in the speed of sound as a function of energy density [26, 237], which differs from purely nucleonic descriptions where the sound speed monotonically increases as illustrated in Fig. 2.7.

At sufficiently high densities in neutron stars, the formation of hyperons through electroweak processes may become energetically favorable [13]. It was frequently argued, however, that the additional degrees of freedom introduced via the hyperons lead to a softening of the equation of state such that heavy neutron stars with $M \sim 2 M_\odot$ cannot be supported against gravitational collapse [56, 238]. This disparity is referred to as the ‘hyperon puzzle’. Introducing repulsive hyperon-nuclear three-body forces is a possible way to inhibit the appearance of hyperons in neutron stars altogether [239–241]. An alternative picture [242] couples baryons (including hyperons) to a density-dependent non-linear scalar field that effectively represents repulsive many-body correlations, such that the required stiffness of the EoS can be maintained even in the presence of hyperons in the neutron star core. A characteristic feature of this model is a sharply dropping speed of sound at the onset density for the appearance of hyperons.

2.3 DEEP LEARNING

In the subfield of machine learning called deep learning, artificial neural networks are trained to model complicated non-linear relationships. This technique has proved very successful in recent years at solving complex problems across various scientific domains. Motivated by the successes of deep learning methods, there has been a growing interest to employ such techniques to study dense matter and neutron stars [143]. One particular class

of neural networks, normalizing flows, can be used to represent complicated probability distributions in high dimensions [145]. In this section we provide a brief overview of the fundamental ideas of deep learning⁴, for more details see, e.g., Refs. [244, 245].

2.3.1 ARTIFICIAL NEURAL NETWORKS

The general objective of deep learning is to fit a non-parametric function f_{Φ} based on the data $\{x_i\}$ to optimize a task-specific cost function. The non-parametric function is represented using artificial neural networks (ANNs), computational structures inspired by the neural architecture of the human brain. ANNs consist of interconnected layers of ‘neurons’. The first layer receives the input data, with each neuron representing a feature or attribute of the input, and the last layer provides the network’s output.




FIG. 2.8. Schematic illustration of a feed-forward neural network with two hidden layers.

In practice, it is advantageous to include intermediate ‘hidden’ layers between the input and output layers. In the most basic architecture of a *feed-forward neural network* (sometimes also called a multilayer perceptron) depicted in Fig. 2.8, each neuron receives inputs from the previous layer, processes them, and produces an output signal that is passed on to the next layer. More specifically, the output of a neuron in the l^{th} layer, $a_i^{(l)}$, is given by a transformation of the neurons in the previous layer, $a_j^{(l-1)}$, involving an activation function $\sigma(z)$ along with adjustable weights $\omega_{ij}^{(l)}$ and biases $b_i^{(l)}$:

$$a_i^{(l)} = \sigma\left(\omega_{ij}^{(l)} a_j^{(l-1)} + b_i^{(l)}\right). \quad (2.30)$$

The activation functions introduce crucial non-linearity to the neural network, allowing it to approximate complex functions. In fact, the universal approximation theorem states that a sufficiently large feed-forward neural network, endowed with appropriate activation functions, can approximate any continuous function on a compact subset of its input space to arbitrary precision [246]. Common activation functions include the sigmoid function, the hyperbolic tangent, and the rectified linear unit (ReLU), $\sigma(z) = \max\{0, z\}$. Beyond feed-forward neural networks, more complex architectures such as convolutional or recurrent neural networks have been designed for specific data modalities like image or sequential input data.

2.3.2 TRAINING

The parameters $\Phi = \{\omega, b\}$ of a neural network are trained to minimize a predefined cost function C (sometimes also called loss function), chosen for the specific problem at hand. Common choices include a mean squared error for regression or a cross-entropy loss for

⁴ Parts of the text in this section have been adapted from Ref. [243].

classification problems. To avoid overfitting, i.e., the neural network learning the noise of the available data at the expense of poor generalization to unseen data, the data set is divided into training, validation and test sets. The training data, $\{x_i\}_{i=0}^N$, is used to train the parameters of the neural network. The total cost function is given by the average over all training samples, $1/N \sum_{i=0}^N C(\Phi, x_i)$.

The parameters Φ are optimized using gradient descent. To compute the necessary gradients of the cost function with respect to the weights, $\nabla_{\Phi} C(\Phi, x_i)$, the chain rule is applied iteratively backwards through the layers of the network, starting from the output layer and moving towards the input layer in a procedure called backpropagation. Using gradient descent, the network parameters are updated with a learning rate α

$$\Phi_{t+1} = \Phi_t - \alpha \frac{1}{N} \sum_{i=0}^N \nabla_{\Phi} C(\Phi, x_i) . \quad (2.31)$$

In scenarios with massive data sets, computing the gradients of the entire training set can become computationally expensive as well as memory intensive. Stochastic gradient descent addresses these issues by computing the gradient using only a small subset of the training data at each iteration t . A commonly used optimization algorithm to make the training more robust and efficient is adaptive moment estimation (Adam), where the learning rate is adaptively adjusted for each parameter [247].

The validation set is used to evaluate the performance of the model during training and to tune hyperparameters, i.e., settings that are not learned from the data but affect the learning process such as the learning rate α or the number of hidden layers in the neural network. The test set serves as a completely independent data set that the model has not seen during training or validation, such that it can be used to provide an unbiased evaluation of the final model after it has been tuned and trained.

2.3.3 NORMALIZING FLOWS

Normalizing flows (NFs) are a class of generative models 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 $p(x)$ [145, 146]. The main idea, illustrated in Fig. 2.9, 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, $f_{\Phi} = f_{\Phi_1} \circ \dots \circ f_{\Phi_N}$, we can easily generate samples x from $p(x)$ by transforming the samples u of $\pi(u)$:

$$u \sim \pi(u) , \quad (2.32)$$

$$x = f_{\Phi}(u) = f_{\Phi_N}(\dots(f_{\Phi_1}(u))) . \quad (2.33)$$

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

$$p(x) = \pi \left(f_{\Phi}^{-1}(x) \right) \left| \det \left(\frac{\partial f_{\Phi}^{-1}}{\partial x} \right) \right| = \pi \left(f_{\Phi_1}^{-1}(\dots(f_{\Phi_N}^{-1}(x))) \right) \prod_{i=1}^N \left| \det \left(\frac{\partial f_{\Phi_i}^{-1}}{\partial x} \right) \right| . \quad (2.34)$$

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 is fast to compute.

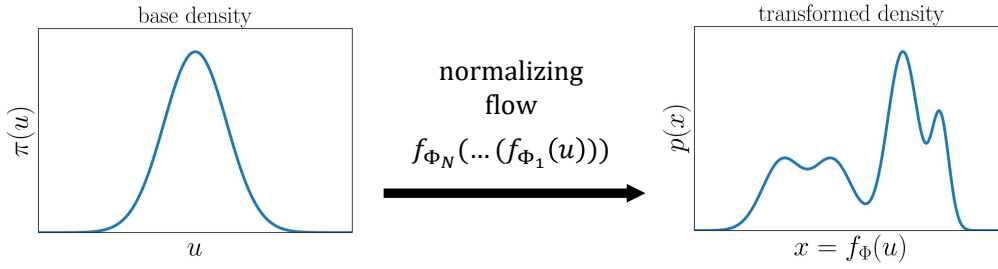


FIG. 2.9. Schematic illustration of the working principle of a normalizing flow in one dimension.

One popular choice are Masked Autoregressive Flows (MAFs) [248], where each dimension x_d of the samples, $x = (x_1, \dots, x_d, \dots, x_D)$, is sequentially transformed conditioned on the previous dimensions, $x_{1:d-1} = x_1, \dots, x_{d-1}$. In autoregressive models the complex probability distribution $p(x)$ is decomposed into a product of one-dimensional conditional densities:

$$p(x) = \prod_{d=1}^D p(x_d | x_{1:d-1}) . \quad (2.35)$$

A common choice for the conditional densities are Gaussian marginals,

$$p(x_d | x_{1:d-1}) = \mathcal{N}(x_d | \mu_d, \sigma_d) , \quad (2.36)$$

whose mean and standard deviations are computed by neural networks based on the previous dimensions:

$$\mu_d = f_{\mu_d}(x_{1:d-1}) , \quad (2.37)$$

$$\sigma_d = f_{\sigma_d}(x_{1:d-1}) . \quad (2.38)$$

We can sample from this conditional distribution by scaling and shifting the samples u_d from a standard Gaussian

$$x_d = \mu_d + \sigma_d u_d , \quad (2.39)$$

$$u_d \sim \mathcal{N}(0, 1) . \quad (2.40)$$

Accordingly, autoregressive models can be interpreted as a transformation f_{Φ_i} of the base distribution $\pi(u) = \mathcal{N}(0, 1)$ [249]. We can stack multiple autoregressive transformations into one deep normalizing flow f_{Φ} . Usually the ordering of dimensions x_1, \dots, x_D is changed for each bijector in the flow. Computing the inverse transformation $f_{\Phi_i}^{-1}$ does not require us to invert the functions f_{μ_d} or f_{σ_d} , since

$$u_d = \frac{x_d - \mu_d}{\sigma_d} . \quad (2.41)$$

This allows us to use non-linear activation functions in the neural networks f_{μ_d} and f_{σ_d} . Due to the autoregressive structure, the Jacobian of $f_{\Phi_i}^{-1}$ is triangular by design and its determinant is easy to compute,

$$\left| \det \left(\frac{\partial f_{\Phi_i}^{-1}}{\partial x} \right) \right| = \prod_{d=1}^D \sigma_d^{-1} . \quad (2.42)$$

With suitably constructed binary masks, we can use a single feed-forward neural network to output all values of μ_d and σ_d , while retaining the autoregressive property [250]. Since the conditional likelihoods in Eq. (2.35) can be computed in parallel, it is possible to recover all values of u_d from all values of x_d with a single pass through the neural network, making MAFs very efficient in computing the probability density $p(x)$. On the other hand, generating new samples, $x \sim p(x)$, is less efficient because all previous $x_{1:d-1}$ must be computed before computing x_d . MAFs have been used to model complex probability distributions in various scientific applications [251–253].

INFERENCE OF THE EQUATION OF STATE

In this chapter, we introduce the methodology necessary to infer the neutron star matter equation of state. As discussed in the previous chapter, the phase structure of strongly interacting matter and thus the EoS at the high densities realized inside neutron stars is still unknown. Therefore, we adopt various general parametrizations of the equation of state, each capable of modeling various phase scenarios depending on their parameters. We then introduce Bayesian inference methods to derive the posterior probabilities of these EoS parameters based on astrophysical neutron star data and theoretical constraints at low and high densities. Bayesian statistics further allow us to compute posterior credible bands for neutron star properties and to conduct hypothesis testing using Bayes factors. Additionally, this chapter reviews the available multimessenger neutron star data and their implementation in the inference procedure. The conventional Bayesian inference of the neutron star matter EoS is carried out based on a set of exterior neutron star properties, such as masses, radii, or tidal deformabilities, previously inferred from observational detector data. At the end of this chapter, we introduce a novel approach that employs recently developed neural simulation-based inference techniques to infer the EoS parameters directly from detector data, without the traditionally required intermediate step via the exterior neutron star properties.

3.1 EQUATION OF STATE PARAMETRIZATIONS

A variety of parametrizations has been introduced to represent the equation of state in neutron stars, among the most prominent ones are piecewise polytropes [254] or spectral representations [255]. As discussed in Sec. 2.2.4, various theories predict different phase structures at high densities including phase transitions or crossovers, which are reflected in the behavior of the speed of sound [136, 256]. In the present analysis we employ three different parametrizations for $c_s^2(\varepsilon)$ inside neutron stars: a skewed Gaussian function, piecewise segmented linear interpolations and a spectral expansion¹. Parametrizations of $P(\varepsilon)$ such as piecewise polytropes can cause unphysical discontinuous effects in the speed of sound. In contrast, the parametrizations employed here are continuous in $c_s^2(\varepsilon)$.

The Gaussian and Segments parametrizations used in our conventional Bayesian inference analysis depend on sets of either six or eight parameters θ . A comparative study using these two different forms will give an impression of possible systematic uncertainties induced by the choice of parametrization. This is accompanied by an improvement in the generality of the Segments parametrization, which increases the number of free parameters to ten. At very low densities, $n \leq 0.5 n_0$, the speed of sound is matched to the neutron star crust

¹ Parts of the text in this section have been adapted from Refs. [96, 97].

modeled by the time-honored Baym–Pethick–Sutherland (BPS) parametrization [257]. The effect of the neutron star crust on observables studied in this work is expected to be small, for the potential impact of the crust on macroscopic neutron star properties see Refs. [258, 259].

3.1.1 GAUSSIAN

Based on Refs. [67, 138] we can represent the squared speed of sound in neutron star matter at zero temperature as a function of energy density by a skewed Gaussian. A logistic function is added such that the parametrization reaches the conformal limit $c_s^2 \rightarrow 1/3$ at asymptotically high energy densities. With $x = \varepsilon/(m_N n_0)$, where m_N is the free nucleon mass, the squared speed of sound is represented as:

$$c_s^2(x, \theta) = a_1 \exp \left[-\frac{1}{2} \frac{(x - a_2)^2}{a_3^2} \right] \left(1 + \operatorname{erf} \left[\frac{a_6}{\sqrt{2}} \frac{x - a_2}{a_3} \right] \right) + \frac{1/3 - a_7}{1 + \exp[-a_5(x - a_4)]} + a_7, \quad (3.43)$$

with $\operatorname{erf}(z) = \frac{2}{\sqrt{\pi}} \int_0^z dt e^{-t^2}$ the conventional error function. The parameter a_7 is determined such that the transition to the neutron star crust is continuous. Hence, six free parameters $\theta = (a_1, \dots, a_6)$ remain. When $c_s^2(x, \theta)$ becomes negative, violating thermodynamic stability of the star, we set $c_s^2 = 0$. In this way the Gaussian parametrization can describe arbitrarily strong phase transitions. The combination of the Gaussian and logistic function can also account for variable crossovers. As argued in Sec. 2.2.4, a local maximum in the speed of sound could indicate a transition from baryonic to quark dynamics [3, 200, 234] or the onset of hyperonic degrees of freedom [242].

3.1.2 SEGMENTS

The Gaussian parametrization assumes a specific functional form of the sound speed inside neutron stars. At the present stage the empirical database is still limited, so that inference procedures can depend sensitively on prior choices including the functional form of the parametrization [68, 76]. For an alternative test, results of broader generality can be produced using a more universal parametrization of the speed of sound based on segment-wise linear interpolations, similar to Refs. [130, 131].

This parametrization is represented by a set of $N + 1$ points $\theta = (c_{s,i}^2, \varepsilon_i)$. The squared speed of sound $c_s^2(\varepsilon, \theta)$ is modeled as a linear interpolation between these points, i.e., for $\varepsilon \in [\varepsilon_i, \varepsilon_{i+1}]$ with $i = 0, \dots, N$ one has:

$$c_s^2(\varepsilon, \theta) = \frac{(\varepsilon_{i+1} - \varepsilon)c_{s,i}^2 + (\varepsilon - \varepsilon_i)c_{s,i+1}^2}{\varepsilon_{i+1} - \varepsilon_i}. \quad (3.44)$$

The $i = 0$ point is the transition point to the neutron star crust $(c_{s,0}^2, \varepsilon_0) = (c_{s,\text{crust}}^2, \varepsilon_{\text{crust}})$. Beyond the last point, $\varepsilon > \varepsilon_N$, the speed of sound becomes constant, $c_s^2(\varepsilon, \theta) = c_{s,N}^2$.

In the ‘basic’ version of the Segments parametrization the last point is chosen such that the conformal limit is reached at very high energy densities $(c_{s,N}^2, \varepsilon_N) = (1/3, 10 \text{ GeV fm}^{-3})$. We have checked that the results do not depend on the specific choice of ε_N as long as its value is large enough. The asymptotic end point at $\varepsilon_N = 10 \text{ GeV fm}^{-3}$ corresponds to a baryon chemical potential of $\mu \sim 2.4 \text{ GeV}$ in the pQCD results from Ref. [49]. The asymptotic conformal limit is approached from below, which is realized by restricting the last speed-of-sound value before the end point to $c_{s,N-1}^2 < 1/3$.

Here we choose $N = 5$, corresponding to a total of eight free parameters and five free segments, which turns out to be enough for a representation of c_s^2 . Ref. [132] uses a similar parametrization based on piecewise segments. There it is found that five segments are sufficient to avoid numerical artifacts, namely that for a larger number of segments the results do not change significantly any more. An equivalent or smaller number of segments is used to interpolate over the full range between ChEFT and pQCD constraints in Refs. [130, 131]. It was found in Ref. [91] that four segments are already sufficient to describe the current astrophysical data and lead to results comparable to those of a non-parametric Gaussian process. We nevertheless study whether increasing the number of segments in the parametrization produces notably different results. The Segments parametrization can also incorporate a variety of phase transitions or crossovers. In contrast to the representation in terms of a skewed Gaussian it can also accommodate possible steep rises as well as plateaus in the speed of sound.

IMPROVED SEGMENTS PARAMETRIZATION

We can increase the generality of the parametrization by not assuming the asymptotic behavior of the sound velocity a priori. Thus, in an ‘improved’ version of the Segments parametrization, the last point ($i = N$) is no longer fixed to reproduce the asymptotic conformal limit, and the second-to-last sound speed is no longer limited to $c_{s,N-1}^2 < 1/3$. The pQCD constraint at asymptotic densities is still systematically implemented through a pQCD likelihood (see Sec. 3.4.5). With the parameters of the last point ($c_{s,N}^2, \epsilon_N$) not assumed to take the value of the conformal limit, the number of free parameters of the Segments parametrization increases to ten and the number of free segments increases by one. Thus, the improved version also allows us to test whether the posterior results are stable with respect to an increase in the number of segments. At the same time, we keep the basic version of the Segments parametrization to allow a more direct comparison with the Gaussian parametrization with the same asymptotic behavior.

A central focus of the present work is the possible occurrence of a first-order phase transition in neutron star matter. In the EoS this corresponds to a jump in energy density, i.e., the appearance of two successive discontinuities in the speed of sound. For instance, in a Maxwell construction a phase coexistence region of constant pressure emerges along a density interval Δn . At the lower end of this interval the sound velocity drops to zero while at the upper end it jumps back to a finite value, see Fig. 2.7. This and similar kinds of scenarios can be represented in the Segments parametrization of Eq. (3.44) when one of the interpolation points reaches a small sound speed, $c_{s,i}^2 \sim 0$, while the two adjacent points remain at finite values. This condition as such is not sufficient to identify a first-order phase transition. However, in combination with a detailed quantitative inspection of $\Delta n/n$ as a measure for the extension of a phase coexistence region that can possibly develop within the posterior credible bands, it serves to set constraints on the appearance of a strong first-order transition. We refer to a first-order transition as ‘strong’ if $\Delta n/n > 1$ (where n is the density at which the coexistence interval ends). In contrast a ‘weakly’ first-order transition has $\Delta n/n$ small compared to one.

3.1.3 SPECTRAL

For the neural simulation-based inference conducted later in this work, we use a spectral expansion of the speed of sound depending only on two parameters to make our results directly comparable to previous approaches [127, 128].

In the spectral approach, a representation of the equation of state is expanded in a series of basis functions. One popular choice is to represent the equation of state in terms of the adiabatic index [255]

$$\Gamma(P) = \frac{\varepsilon(P) + P \frac{dP}{d\varepsilon}}{P}. \quad (3.45)$$

If $\Gamma(P)$ is non-negative, the EoS remains thermodynamically stable, but can become acausal. Therefore, we choose an alternative representation of the equation of state in terms of a function of the squared speed of sound c_s^2 [260]

$$\Upsilon(P) = \frac{1 - c_s^2}{c_s^2}, \quad (3.46)$$

which, if larger than zero, is causal by construction. $\Upsilon(P)$ can be parametrized via an expansion in the pressure. However, since the numerical solution of the TOV equations is more convenient in terms of the effective enthalpy h , as explained in Appendix A, it is preferable to expand the sound speed function in this quantity instead:

$$\Upsilon(h) = \exp \left[\sum_k \lambda_k \Phi_k(h) \right]. \quad (3.47)$$

Here, $\Phi_k(h)$ can be any complete set of basis functions. The most commonly used choice is $\Phi_k = [\log(h/h_0)]^k$, where h_0 is the matching point to the low-density crust. For given expansion coefficients λ_k , pressure and energy density can be determined via enthalpy integrations over the sound speed function. A spectral expansion with only two coefficients is already sufficient to describe a variety of different equations of state [254] with error residuals below 10% [127, 255]. Note, however, that spectral parametrizations are less well suited to describe extended first-order phase transitions [260].

3.2 BAYESIAN INFERENCE

Making use of a set of available neutron star data, we aim to constrain the free parameters θ of the parametrizations described in the previous section. For that purpose we use Bayesian inference, similar to Refs. [2, 63–107], and follow Refs. [139, 140]². More specifically, Bayes' theorem allows us to compute the probability distribution for the parameters θ based on the available data. To include the information from astrophysical observations in the inference procedure we need to employ kernel density estimation techniques. From the probability distribution of the parameters we can compute marginal credible bands to constrain the equation of state. Moreover, using Bayesian statistics we can compare the evidence for two competing hypotheses using Bayes factors.

3.2.1 BAYES' THEOREM

For given data \mathcal{D} , the *posterior probability distribution* for the parameters θ can be computed using Bayes' theorem:

$$p(\theta|\mathcal{D}) = \frac{p(\mathcal{D}|\theta) p(\theta)}{p(\mathcal{D})}. \quad (3.48)$$

² Parts of the text in this section have been adapted from Refs. [96, 97].

For the sake of readability, we have neglected here the conditional dependence on the model \mathcal{M} which includes all implicit assumptions, such as the parametrization choice, etc., and is sometimes made explicit in the literature. The probability distribution $p(\theta)$ for the EoS parameters is denoted the *prior* and given by the chosen parameter distributions for each parametrization. The probability $p(\mathcal{D}|\theta)$ for the data \mathcal{D} to occur, given the parameters θ , is usually referred to as the *likelihood*.

To compute the likelihood of the neutron star data for a given EoS characterized by parameters θ , we have to consider all possible neutron stars supported by this particular EoS. This means we have to marginalize over all central pressures P_c (or, alternatively, central energy densities ε_c) possibly realized inside neutron stars,

$$p(\mathcal{D}|\theta) = \int dP_c p(\mathcal{D}|\theta, P_c)p(P_c(\theta)) . \quad (3.49)$$

This is weighted by a prior probability distribution for the central pressures, $p(P_c(\theta))$, which allows to take the distribution of the neutron star population into account. The prior distribution of the central pressures depends on θ because the maximum central pressure leading to a stable solution, $P_{c,\max}$, is different for each set of parameters. By numerically solving the coupled system of TOV equations and the equations for the tidal deformability in Sec. 2.1, a set of parameters θ and central pressures P_c is deterministically linked to a mass-radius relation (M, R) and tidal deformabilities Λ . Therefore, we can write:

$$p(\mathcal{D}|\theta, P_c) = p(\mathcal{D}|M, R, \Lambda) . \quad (3.50)$$

Similarly, with the radius and tidal deformability available as functions of the mass, $R(M, \theta)$ and $\Lambda(M, \theta)$ after solving the differential equations in Sec. 2.1 for given θ , the central pressure prior $p(P_c(\theta))$ defines a mass prior $p(M(\theta))$. The likelihood $p(\mathcal{D}|M, R, \Lambda)$ is usually analytically unavailable, therefore we have to resort to a two-step inference approach [67, 139], i.e., for computational feasibility we have to assume that we can use the posterior distributions from the analyses of neutron star observables as likelihoods for exterior neutron star parameters:

$$p(\mathcal{D}|M, R, \Lambda) \propto p(M, R, \Lambda|\mathcal{D}) . \quad (3.51)$$

This is valid if the prior of (M, R, Λ) used in the inference analyses of the observational data is sufficiently uninformative, which is the case for the majority of the data analyzed in this work as argued in Ref. [70]. Later in Sec. 3.5, we introduce a single-step approach for the inference directly based on the data \mathcal{D} using deep learning techniques. For now, in the conventional two-step approach, the likelihood can be evaluated based on the inferred posterior constraints as

$$p(\mathcal{D}|\theta) \propto \int dM p(M, R, \Lambda|\mathcal{D})p(M(\theta)) , \quad (3.52)$$

where the radius and tidal deformability are given by the functions $R(M, \theta)$ and $\Lambda(M, \theta)$ for given θ . We introduce the parameter priors, $p(\theta)$, as well as the mass prior, $p(M(\theta))$, used in our Bayesian analysis in Sec. 3.3 and give an overview over the currently available neutron star data in Sec. 3.4.

The probability distribution $p(\mathcal{D})$ in the denominator of Eq. (3.48) is usually referred to as the *evidence* or marginal likelihood. It is determined by the normalization of the posterior:

$$p(\mathcal{D}) = \int d\theta p(\mathcal{D}|\theta)p(\theta) . \quad (3.53)$$

Depending on the number of parameters this may be a high-dimensional integral which can be difficult to solve numerically. Therefore, in Bayesian inference sampling algorithms such as Markov Chain Monte Carlo (MCMC) or nested sampling are frequently employed. For a sufficiently low-dimensional parameter space, samples from the prior, $p(\theta)$, weighted with the likelihood, $p(\mathcal{D}|\theta)$, yield the posterior probability distribution up to a multiplicative constant. In this case it needs to be checked whether the number of samples is large enough such that sufficient probability mass of the posterior has been covered. Similar importance sampling techniques have frequently been employed in Bayesian analyses of neutron stars (see for example Refs. [91, 261]).

3.2.2 KERNEL DENSITY ESTIMATION

The posterior distributions inferred from the neutron star observations, $p(M, R, \Lambda|\mathcal{D})$, are usually available only as samples. We can determine the underlying probability distributions using *kernel density estimation*, which is a non-parametric method to determine the probability density function of a given data set. Assume a set of N points, (x_1, \dots, x_N) , which are independent and identically distributed according to the unknown density function $p(x)$. The kernel density estimator (KDE) of this underlying density function is

$$\text{KDE}(x) = \frac{1}{Nh} \sum_{i=1}^N K\left(\frac{x - x_i}{h}\right), \quad (3.54)$$

where h is a smoothing parameter called the bandwidth and K is a kernel function. This kernel function must integrate to one and be symmetric and non-negative. There is a range of possible kernel functions, e.g., uniform, linear or exponential and here we use a normalized Gaussian kernel:

$$K(z) = \frac{1}{\sqrt{2\pi}} \exp(-z^2/2). \quad (3.55)$$

The choice of the bandwidth h is done such that a proper balance is achieved between maintaining important features in the density function and smoothing over irrelevant fine structure in the estimator. To find an appropriate value for h we use the rule of thumb developed by Silverman [262]. The above approach can be straightforwardly generalized to the case with data on a higher dimensional space.

3.2.3 CREDIBLE BANDS

To transform the posterior distribution on the parameters θ to the EoS space, we follow Ref. [67] in discretizing energy densities on a grid $\{\varepsilon_i\}$. For each posterior sample the pressure is determined at each discrete energy density $P(\varepsilon_i, \theta)$, up to the maximum central energy density $\varepsilon_{c,\text{max}}$, corresponding to the endpoint M_{max} of the mass-radius relation. In this way we obtain the marginal posterior distribution for the pressure $p(P|\varepsilon_i, \mathcal{D})$ at each energy density. We can then determine the *credible interval* $[a, b]$ at the levels $\alpha = 68\%$ or 95% as

$$\alpha = \int_a^b dP \ p(P|\varepsilon_i, \mathcal{D}). \quad (3.56)$$

In principle, there is an infinite amount of possible intervals $[a, b]$ for a given α . Here, we choose to use highest density credible intervals, which correspond to the narrowest interval leading to the probability α . Combining the credible intervals at each ε_i gives a

posterior credible band for $P(\varepsilon)$. Similarly we can find credible bands for $c_s^2(\varepsilon)$, $R(M)$, $\Lambda(M)$, etc.. In contrast, displaying neutron star properties such as the EoS, $P(\varepsilon)$, via a two-dimensional credible region depends on the chosen prior in ε , so that different prior choices can lead to different results. Hence in the literature, with few exceptions, the procedure in terms of credible bands is favored [67–70, 72, 74, 88]. Note that each EoS characterized by parameters θ is only used up to the maximum central energy density $\varepsilon_{c,\max}$, the one corresponding to the maximum TOV mass M_{\max} supported by that EoS. The latter statement applies unless the EoS leads to more than one stable branch in the mass-radius relation. At higher energy densities (or masses), the credible intervals computed via Eq. (3.56) are determined on the basis of correspondingly fewer equations of state. This loss of expressive power at higher energy densities and masses is not reflected in the credible bands.

3.2.4 BAYES FACTORS

With *Bayes factors* one can compare two competing hypotheses H_0 and H_1 and quantify the evidence for one hypothesis over the other. Given a data set \mathcal{D} , the Bayes factor is defined as the quotient of the marginal likelihoods:

$$\mathcal{B}_{H_0}^{H_1} = \frac{p(\mathcal{D}|H_1)}{p(\mathcal{D}|H_0)}. \quad (3.57)$$

The marginal likelihood for a general hypothesis H is given by the integral over all parameter sets that support this hypothesis, $\theta \in \Theta_H$:

$$p(\mathcal{D}|H) = \int_{\theta \in \Theta_H} d\theta p(\mathcal{D}|\theta)p(\theta), \quad (3.58)$$

which we can approximate using Monte-Carlo integration,

$$p(\mathcal{D}|H) \approx \frac{1}{N_{\theta \in \Theta_H}} \sum_{\theta \in \Theta_H} p(\mathcal{D}|\theta), \quad (3.59)$$

if we sum over parameter sets which are sampled from the prior, $\theta \sim p(\theta)$. $N_{\theta \in \Theta_H}$ denotes the number of parameter sets that support the hypothesis H . This holds only if the number of samples is large enough to approximate the integration over the multidimensional parameter space by a sum. To interpret the resulting Bayes factors we use the established evidence classification of Ref. [263, 264] listed in Table 3.1.

3.3 PRIORS

To compute the posterior probability distribution (3.48), prior distributions for the EoS parameters $p(\theta)$ and for the mass distribution of neutron stars $p(M(\theta))$ must be chosen, which will be specified in this section³. The prior for the spectral parametrization follows previous work [127]. In addition, we also introduce a test prior for a restrictive scenario with monotonically rising sound velocity to investigate a purely hadronic composition of neutron star matter.

³ Parts of the text in this section have been adapted from Refs. [96, 97, 243].

$B_{H_0}^{H_1}$	Interpretation
> 100	Extreme evidence for H_1
30 - 100	Very strong evidence for H_1
10 - 30	Strong evidence for H_1
3 - 10	Moderate evidence for H_1
1 - 3	Anecdotal evidence for H_1
1	No evidence
1/3 - 1	Anecdotal evidence for H_0
1/10 - 1/3	Moderate evidence for H_0
1/30 - 1/10	Strong evidence for H_0
1/100 - 1/30	Very strong evidence for H_0
$< 1/100$	Extreme evidence for H_0

TABLE 3.1. Interpretation of Bayes factors for comparing the evidence for hypotheses H_0 and H_1 according to the evidence classification in Ref. [263] with the updated terminology of Ref. [264].

3.3.1 PARAMETER PRIORS

In order to ensure maximum generality of the results we choose very broad parameter ranges for all parametrizations, covering most of the speed-of-sound space. However, for the Gaussian and Segments parametrization with fixed asymptotic behavior we follow previous work [68–70] to discard parameter sets that lead to multiple disconnected stable mass-radius relations. To obtain even more general results, we choose no longer to reject parameter sets that lead to disconnected branches for the improved version of the Segments parametrization, so that this prior set includes the possibility of twin-star scenarios and even cases with more than two disconnected branches in the mass-radius relation.

GAUSSIAN

The six free parameters of the Gaussian parametrization (referred to in the following as version G) are sampled from uniform intervals listed in Table 3.2. These parameter ranges were chosen guided by previous studies [67, 138]. The resulting functions cover the speed-of-sound space sufficiently well. Only those combinations of parameters are kept that lead to causal EoS, i.e., $c_s^2(\varepsilon) < 1$ for all energy densities. In our default version G the asymptotic conformal limit, $c_s^2 = 1/3$, is approached from below as in standard pQCD [197], implying that the derivative of the speed of sound must be positive, $\partial c_s^2 / \partial \varepsilon > 0$, at very high energy densities. In practice this onset of asymptotic behavior is imposed at three different values, $\varepsilon = 4, 8$ and 16 GeV fm^{-3} . We have checked that this specific choice does not affect the inference results as long as these energy densities are sufficiently large.

SEGMENTS

The ‘basic’ version of the segment-wise parametrization with fixed asymptotic behavior (referred to in the following as version S) depends on four speeds of sound and energy densities ($c_{s,i}^2, \varepsilon_i$). The number of parameters increases to ten for the more general variant (which we will refer to as version S’). In both cases the energy densities are sampled logarithmically from $\varepsilon_i \in [\varepsilon_{\text{crust}}, 4 \text{ GeV fm}^{-3}]$, where $\varepsilon_{\text{crust}}$ refers to the endpoint of the neutron star crust.

Parameter	Range
a_1	[0.2, 3]
a_2	[0.5, 12]
a_3/a_2	[0.05, 10]
a_4	[0.1, 15]
a_5	[0.1, 5]
a_6	[-15, 15]

TABLE 3.2. Prior ranges for the six parameters of the Gaussian parametrization of the speed of sound inside neutron stars given in Eq. (3.43).

With this sampling the large multitude of EoSs in the prior is represented, on average, by three to four free segments. The speed-of-sound values are collected from logarithmic intervals $c_{s,i}^2 \in [0, 1]$, so they are causal by construction and at the same time open to the possible occurrence of phase transitions. Note that the last speed-of-sound value before the end point is restricted to $c_{s,N-1}^2 < 1/3$ in the S version so that the asymptotic conformal limit is approached from below. With at least two more parameters and a more general functional form, the Segments parametrization allows, in principle, to describe more complex structures compared to the Gaussian parametrization. As a stability test we have checked that shifting the upper limit of the logarithmic interval downward from its value $\varepsilon_{i,\max} = 4 \text{ GeV fm}^{-3}$ induces small changes in the prior but does not affect the final posterior results.

The prior credible bands for both G and S parametrizations are depicted in Fig. 3.1. The bands are very broad in both interior and exterior parameter spaces. Because the ChEFT constraint is employed as likelihood and hence not present in the priors (see Sec. 3.4.5), there is prior support for rapidly increasing speeds of sound at low densities, leading to large neutron star radii. Hence the prior credible bands have strong weights both at small sound speeds and large radii.

The parameter ranges are chosen to minimize any possible restrictions, such that the posterior distribution has maximum freedom to be governed by the empirical data. The prior probability distributions of the G and S versions differ because of the different functional forms and chosen parameter ranges. This permits an assessment of the impact of different prior choices on the inference results. If the results for versions G and S turn out to be very similar, we can conclude that the inference procedure is robust against variations in the functional form of the prior. The prior probability distribution for $c_s^2(\varepsilon)$ in the S' version is displayed in Fig. 3.2. The prior credible bands looks similar to those of the basic version of the Segments parametrization in Fig. 3.1, although they extend to slightly smaller values. By comparing the inference outputs for the S and S' versions, we can access the impact of increasing the number of free segments and changing the asymptotic sound speed behavior on the posterior results.

All priors at the 95% level support very small speeds of sound, $c_s^2 \lesssim 0.05$. In fact the 68% credible band of the G version reaches down to $c_s^2 = 0$. Accordingly, every fourth EoS in the Gaussian parametrization potentially has a first-order phase transition in the sense that the minimum speed of sound becomes smaller than $c_{s,\min}^2 \leq 0.1$, whereas 15% to 20% of the EoSs in the priors of both versions of the Segments parametrization feature such a phase transition. In contrast, each EoS in the (later determined) posterior credible bands

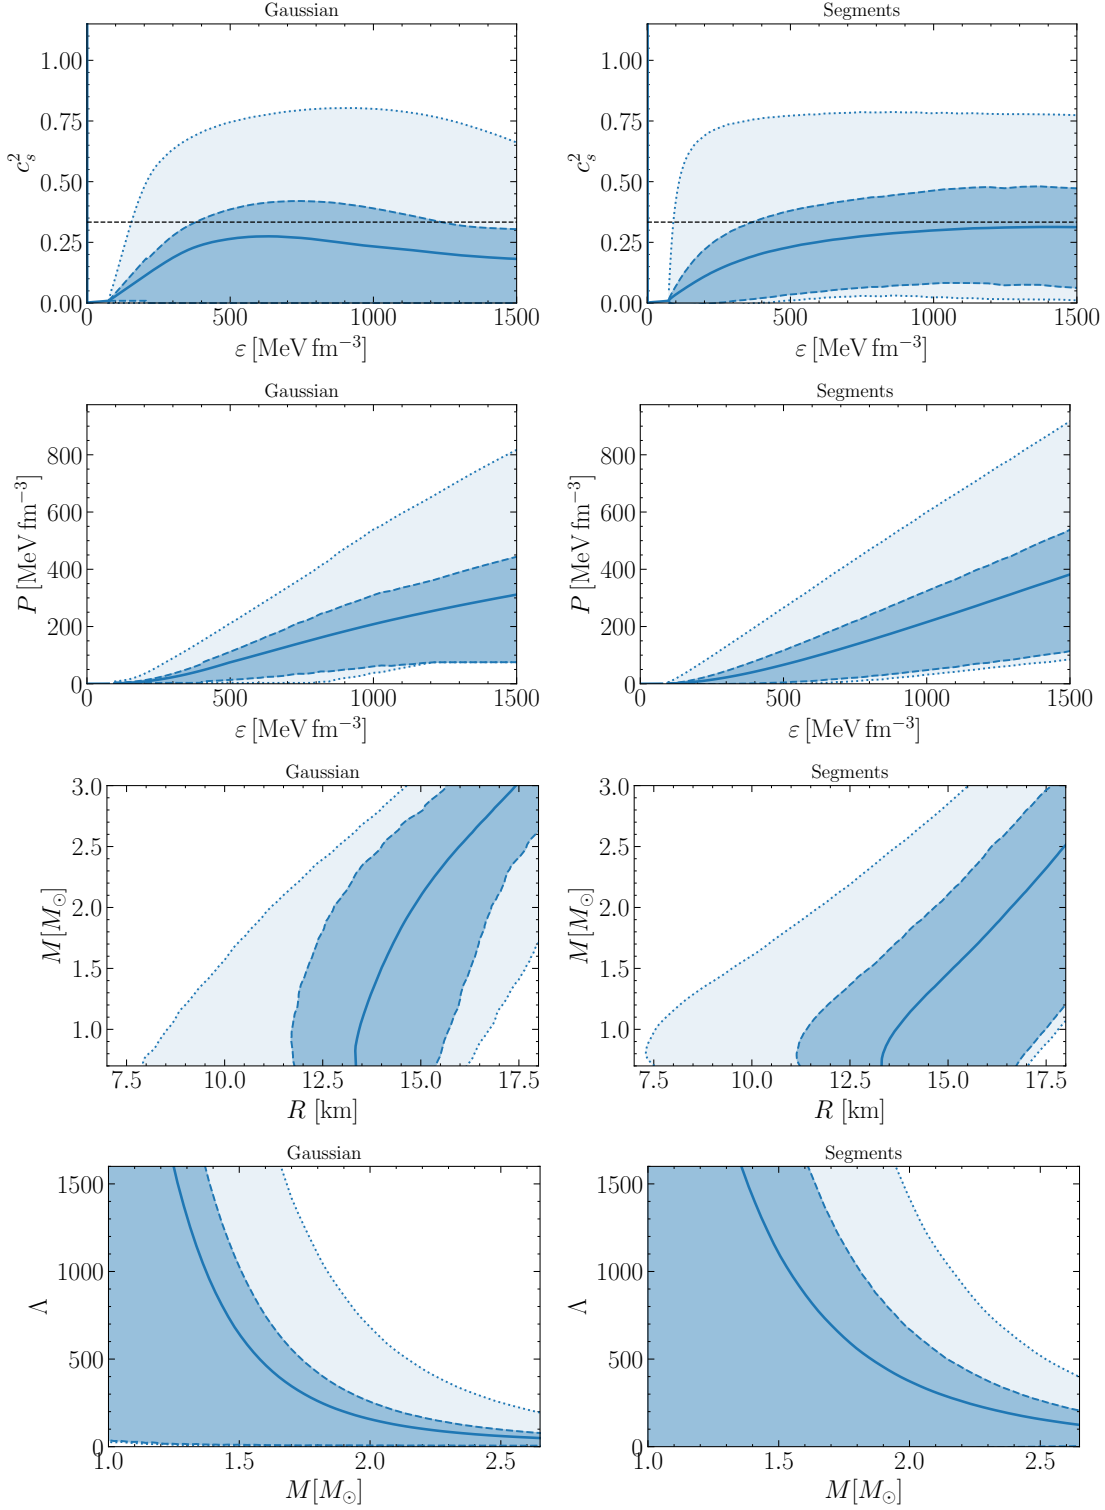


FIG. 3.1. Marginal prior probability distributions at the 95% and 68% level of the squared speed of sound c_s^2 and pressure P as a function of energy density ϵ for the Gaussian (left) and the basic version of the Segments parametrization (right). Also shown is the prior for the mass-radius relation and the tidal deformability, Λ , as a function of neutron star mass M in units of the solar mass M_\odot . At each ϵ or M , there exist 95% and 68% prior credible intervals for $c_s^2(\epsilon)$, $P(\epsilon)$ or $R(M)$, $\Lambda(M)$. These intervals are connected to obtain the prior credible bands. Similarly, the medians of the prior probability distributions at each ϵ or M are connected (solid lines). For the speed of sound the dashed black line indicates the value of the conformal bound. Figure taken from Ref. [96].




FIG. 3.2. Similar to Fig. 3.1: median and prior credible bands for the squared speed of sound, c_s^2 , as a function of energy density, ε , are displayed for the improved version of the Segments parametrization without assumed asymptotic behavior and two more free parameters. Figure taken from Ref. [97].

is constrained by astrophysical data and thus limited by its emerging maximum central energy density, $\varepsilon_{c,\max}$. The mass-radius trajectory deduced from each given EoS, with or without a phase transition, terminates at this point. An EoS's mass-radius sequence normally ends after a first-order phase transition. As a consequence small sound speeds appear with lower weight in the posteriors than in the priors.

3.3.2 MASS PRIOR

Each EoS characterized by a parameter set θ supports neutron star masses between some minimum mass M_{\min} and a respective maximum mass $M_{\max}(\theta)$. We follow Ref. [72] and assume a flat prior distribution between M_{\min} and $M_{\max}(\theta)$:

$$p(M(\theta)) = \begin{cases} \frac{1}{M_{\max}(\theta) - M_{\min}} & \text{if } M \in [M_{\min}, M_{\max}(\theta)] , \\ 0 & \text{else ,} \end{cases} \quad (3.60)$$

with a minimum mass of $M_{\min} = 0.5 M_{\odot}$. This introduces an Occam factor penalizing EoSs that involve extreme masses beyond those supported by the astrophysical data.

The employment of a uniform mass prior differs from some previous works [67, 96, 140] where a central pressure prior was used instead. In Ref. [86] the authors found just a marginal distinction between a flat mass prior and a central pressures prior. However, the uniform mass prior has the advantage that it permits a more direct comparison of our results with other recent works (for example Refs. [72–74, 82, 84–86, 92–94]). Furthermore, for future developments it offers the possibility of easily incorporating the mass population of neutron stars. When the number of available data increases, failing to account for the correct population model may cause a bias in the resulting posterior distribution [265], such that the neutron star population may have to be inferred together with the posterior [266]. Note that this would also allow to estimate the maximum supported mass M_{\max} from the mass distribution of neutron stars [102].

3.3.3 SPECTRAL PARAMETRIZATION

Following Ref. [127], to generate samples for our simulation-based inference analysis, we start from the relativistic mean-field model GM1L [267], which describes neutron star

matter as a collection of baryons interacting through effective meson exchanges. The coupling constants of the model are fitted to reproduce empirical saturation properties of nuclear matter (see for example Ref. [243] for more details). The sound velocity of the GM1L model is represented by a second-order spectral expansion using the process outlined in Sec. 3.1.3. Thus, the full EoS can be succinctly described by the expansion coefficients, which we refer to as λ_1 and λ_2 . The coefficients describing the GM1L EoS are then used to generate many EoS scenarios by varying the parameters around the original values. 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 $\lambda_2 \in [-2.05, -1.85]$. The TOV equations are numerically solved to produce mass-radius relations which are sampled to generate the mass M and radius R for ~ 100 stars per EoS sample. For the sampling of stars we follow Ref. [127] to choose a log-uniform distribution of central enthalpies [166] for the boundary condition to solve the stellar structure equations in Appendix A. In contrast to the uniform mass distribution in Eq. (3.60), this leads to a higher weighting of larger masses close to the maximum supported mass in the distribution of stars for each EoS.

3.3.4 MONOTONICALLY RISING SPEED OF SOUND

The previous, general choice of priors is open, in principle, to possible phase transitions in the EoS if the data suggest such an option. At the same time we also, additionally, wish to investigate a more restrictive case, namely the hypothesis that neutron star matter is composed of conventional hadronic (nucleon and meson) degrees of freedom, with spontaneously broken chiral symmetry intact and no complex phase structure. A successful historical example of this kind is the Akmal–Pandharipande–Ravenhall (APR) equation of state [268]. Another example, which we will discuss in more detail in Sec. 4.2.3, is a model based on chiral nucleon-meson field theory treated non-perturbatively using functional renormalization group methods [219, 220].

With the aim of studying whether such a picture is compatible with the empirical database, the additional prior assumption is implemented that neutron star matter displays no phase transition or crossover up to a given transition density n_{tr} . This is equivalent to the speed of sound rising monotonically up to n_{tr} :

$$\frac{\partial c_s^2}{\partial \varepsilon} > 0 \quad \text{for } n < n_{\text{tr}}. \quad (3.61)$$

For densities $n > n_{\text{tr}}$, the system is allowed any freedom to undergo transitions or changes of degrees freedom. In practice, based on the findings in [221], we vary the transition density in the range $n_{\text{tr}} = 2 - 5 n_0$.

3.4 DATA & LIKELIHOODS

In order to compute the posterior probability distribution in Eq. (3.48) it is necessary to compute likelihoods for the different types of data. In this section a brief account is given of the sets of empirical data used in the inference procedure⁴. This includes measurements of neutron star masses, radii and tidal deformabilities, as well low-density constraints from nuclear theory and high-density constraints from perturbative QCD. For detailed reviews on the current (and upcoming) multimessenger measurements of neutron stars see Refs. [2,

⁴ Parts of the text in this section have been adapted from Refs. [96, 97].

160, 269], as well as Ref. [103] for a detailed examination of the observational impacts on the neutron star matter EoS.

3.4.1 MASS MEASUREMENTS

SHAPIRO TIME DELAY MEASUREMENTS

The most precise mass measurements are available for neutron stars in binary systems with companion stars. Using Kepler's third law, we can derive the binary mass function, which establishes a lower bound for the mass M_1 of the neutron star:

$$f(M_1, M_2) = \frac{M_1^3 \sin^3 i}{(M_1 + M_2)^2} = \frac{K_2^3 P_{\text{orb}}}{2\pi G_N} (1 - e^2)^{3/2}, \quad (3.62)$$

where M_2 represents the mass of the companion and i is the inclination angle of the orbital plane relative to our line of sight, i.e., an inclination of $i \sim 90^\circ$ indicates an edge-on viewing perspective on the binary orbit.

The right-hand side of the expression in Eq. (3.62) consists of observable quantities such as the companion's radial velocity K_2 , the orbital period P_{orb} , and the orbital eccentricity e . For a pulsar companion these parameters can be measured through radio timing, or alternatively, by optical spectroscopy in the case of white dwarf or non-compact stellar companions [160]. In order to overcome the degeneracy of the neutron star's mass with the companion mass and the inclination angle on the left-hand side of the binary mass function, at least two additional properties of the system must be determined.

This can be achieved by measuring relativistic effects beyond the predictions of Newtonian gravity that influence the pulsar signal. One such effect is the *Shapiro time delay*, a periodic variation in the arrival time of signals caused by the photons passing through the gravitational field of the companion star. The magnitude of this delay depends on the mass of the companion star, while its shape depends on the system's inclination [270, 271]. Consequently, this phenomenon is most pronounced in systems that are observed edge-on. Once these parameters have been determined, the neutron star mass can be extracted with high precision.

The most interesting measurements are those of the heaviest neutron stars as this sets a lower limit on the maximum mass that the neutron star EoS has to support. Through Shapiro time delay measurements several neutron stars with masses around twice the solar mass were established, namely PSR J1614-2230 [32-34], PSR J0348+0432 [35] and PSR J0740+6620 [36, 37], with masses evaluated at the 68% level:

$$\text{PSR J1614-2230} \quad M = 1.908 \pm 0.016 M_\odot, \quad (3.63)$$

$$\text{PSR J0348+0432} \quad M = 2.01 \pm 0.04 M_\odot, \quad (3.64)$$

$$\text{PSR J0740+6620} \quad M = 2.08 \pm 0.07 M_\odot. \quad (3.65)$$

To compute the respective likelihoods we follow previous analyses [69, 70, 78] and assume that the mass measurements based on the Shapiro time delay are distributed as Gaussians, $\mathcal{N}(M, \langle M \rangle, \sigma_M) = 1/\sqrt{2\pi\sigma^2} \exp[-1/2 (M - \langle M \rangle)^2/\sigma_M^2]$ with mean values $\langle M \rangle$ and standard deviations σ_M . For a given set of parameters θ the solution of the TOV Eqs. (2.8 - 2.9) yields the maximum supported mass $M_{\text{max}}(\theta)$ for this respective EoS. Then, the likelihood is equivalent to the integral over the Gaussian probability distribution, weighted

with the mass prior in Eq. (3.60):

$$\begin{aligned}
 p(\mathcal{D}_{\text{Mass}}|M(\theta)) &= \int_{M_{\text{min}}}^{M_{\text{max}}(\theta)} dM \mathcal{N}(M, \langle M \rangle, \sigma_M) p(M(\theta)) \\
 &\approx \frac{1}{2} \left[1 + \text{erf} \left(\frac{M_{\text{max}}(\theta) - \langle M \rangle}{\sqrt{2}\sigma_M} \right) \right] p(M(\theta)) .
 \end{aligned} \tag{3.66}$$

The total likelihood for all Shapiro time delay measurements is given by the product of the individual likelihoods.

BLACK WIDOW PULSARS

Optical light curve modeling can be used to determine the mass of so-called ‘spiders’, a class of rapidly rotating pulsars with very low-mass companions. Such binary systems are typically formed when a star evolves into a red giant and its material is accreted by an accompanying neutron star. The neutron star’s rotation is sped up due to the accretion process, and the accelerated pulsar emits a stream of high-energy particles that heat and gradually evaporate the donor star. Pulsars paired with companions weighing around $0.01 M_{\odot}$ are classified ‘black widows’, whereas those with companion masses of the order of $0.1 M_{\odot}$ are termed ‘redbacks’ [160].

In these systems, the tidally locked companion is deformed and heated on the side facing the pulsar. This affects the companion’s optical light curve depending on the inclination of the system. Once the inclination has been determined from these effects, the mass of the neutron star can be derived from the radial velocities of the pulsar and the companion by dividing the binary mass functions (3.62) of the two stars. Note that it is necessary to adjust the spectroscopic measurements of the companion used to determine its radial velocity to account for the asymmetric heating that shifts the center of light away from the geometric center of the star.

Using this technique the heaviest neutron star observed so far was recently reported in Ref. [44]. The black-widow pulsar PSR J0952-0607 is estimated to have a total mass $M = 2.35 \pm 0.17 M_{\odot}$, significantly larger than previously observed masses based on Shapiro time delays, but with a relatively large uncertainty. For this object a simple direct heating model provided the best fit and the results were robust to the inclusion of more complex heating effects. With a rotational frequency $\nu = 709$ Hz, PSR J0952-0607 is also among the fastest-spinning pulsars. Therefore, rotation corrections have to be considered as they can effectively increase the maximum mass that a given EoS can support. We use the radius-dependent rotation adjustment provided in Ref. [272], where the authors find approximately universal relations between stationary and rotating masses and radii. They fit empirical formulas, independent of the EoS, to derive the mass and radius corrections induced by a rotation with frequency ν . We have collected the results of their fit in Appendix B. The resulting radius-dependent rotation correction of the PSR J0952-0607 mass is illustrated in Fig. 3.3. This adjustment is quite strongly radius-dependent, with a significant difference between rotating and non-rotating mass at larger radii. The uncertainty in the correction as reported in Ref. [272] is smaller than the uncertainty in the heavy-mass measurement itself. At a neutron star radius of $R = 12$ km the rotating mass $M_{\text{rot}} = 2.35 M_{\odot}$ decreases by 3% to an equivalent non-rotating mass of $M \simeq 2.28 M_{\odot}$.

In order to compute the black-widow likelihood for a given set of parameters θ , we first determine the TOV maximum non-rotating mass $M_{\text{max}}(\theta)$ as well as the corresponding radius $R(M_{\text{max}}, \theta)$. With the formulas derived in Ref. [272] (see Appendix B) we can




FIG. 3.3. PSR J0952-0607 mass measurement with 68% uncertainty [44] compared to the static case (without the $\nu \simeq 709$ Hz rotation) computed using the prescription in Ref. [272] that is summarized in Appendix B. Figure taken from Ref. [97].

then compute the rotation-adjusted maximum mass $M_{\text{rot,max}}(\theta)$. Assuming that the mass distribution of PSR J0952-0607 is Gaussian, its likelihood can be computed in a way similar to the Shapiro time delay likelihood in Eq. (3.66), with the non-rotating maximum mass replaced by $M_{\text{rot,max}}(\theta)$.

In addition to PSR J0952-0607 there exist further massive spiders, such as PSR J2215+5135, a redback pulsar with a reported mass of $M = 2.15 \pm 0.10 M_{\odot}$ [273]. However, like other spider pulsars, this object required a more complex heating model which introduces a larger systematic uncertainty [44, 273–275]. It is therefore not included in the present analysis.

3.4.2 MASS-RADIUS MEASUREMENTS

PULSE PROFILE MODELING

The mass and radius of neutron stars can be extracted by modeling the thermal X-ray emissions from hot spots that form on their surfaces. For rotation powered, non-accreting pulsars, these hot spots originate from return currents in the pulsar magnetosphere that deposit energy in the surface layers at the magnetic polar caps [276]. The emitted soft X-rays have very stable pulse profiles that vary periodically with the rotation of the star. As these photons traverse the curved spacetime surrounding the neutron star, they are affected by relativistic effects such as gravitational light bending. Due to this effect hot spots remain visible even when they face away from the observer, so that the amplitude of the pulsations is modified by the star’s compactness M/R [277, 278]. In addition, the rotation of the star induces a Doppler effect, leading to an asymmetry in the rise and fall times of the pulsations. Neutron stars with larger radii have higher surface velocities for a given angular velocity, leading to more pronounced Doppler boosts [160]. In order to put tight constraints on the radius, it is therefore advantageous to investigate rapidly rotating sources where the Doppler effects are more prominent.

Using a model of the hot spots and the neutron star atmosphere, Bayesian posterior distributions for the mass and radius can be inferred from X-ray profiles measured by the Neutron Star Interior Composition Explorer (NICER) on board the International Space Station. So far two neutron stars were measured and subsequently analyzed by two independent groups.

Here, we use the results of Riley *et al.* [40, 41] for masses and radii at the 68% level:

$$\begin{aligned} \text{PSR J0030+0451} \quad R &= 12.71_{-1.19}^{+1.14} \text{ km} , \\ M &= 1.34_{-0.16}^{+0.15} M_{\odot} , \end{aligned} \quad (3.67)$$

$$\begin{aligned} \text{PSR J0740+6620} \quad R &= 12.39_{-0.98}^{+1.30} \text{ km} , \\ M &= 2.072_{-0.066}^{+0.067} M_{\odot} . \end{aligned} \quad (3.68)$$

The complementary NICER data analysis by Miller *et al.* found slightly different results [42, 43]:

$$\begin{aligned} \text{PSR J0030+0451} \quad R &= 13.02_{-1.06}^{+1.24} \text{ km} , \\ M &= 1.44_{-0.14}^{+0.15} M_{\odot} , \end{aligned} \quad (3.69)$$

$$\begin{aligned} \text{PSR J0740+6620} \quad R &= 13.7_{-1.5}^{+2.6} \text{ km} , \\ M &= 2.08 \pm 0.07 M_{\odot} , \end{aligned} \quad (3.70)$$

but they are consistent within uncertainties. Based on previous studies we expect the choice between these two analyses to have only little influence on the conclusions [74]. The posterior probabilities for the mass and radius data are available as samples and we can approximate the underlying distributions using kernel density estimation (KDE), see Sec. 3.2.2. For a given EoS characterized by parameters θ , we solve the TOV equations to obtain $R(M, \theta)$. The likelihood is then computed as the mass integral over the KDE evaluated at the radius given by the mass-radius relation:

$$p(\mathcal{D}_{\text{Mass-radius}} | (M, R)(\theta)) = \int_{M_{\min}}^{M_{\max}(\theta)} dM \text{ KDE}(M, R(M, \theta)) p(M(\theta)) . \quad (3.71)$$

This is again weighted with the mass prior in Eq. (3.60). In the case of PSR J0740+6620, the analysis of the NICER measurement includes the Shapiro delay result in Eq. (3.65). In this case we do not include the Shapiro mass data in the total likelihood to avoid double counting.

Recently, a new analysis of PSR J0030+0451 that employed an updated instrument response and additionally included XMM-Newton data was carried out in Ref. [279]. In this analysis, the preferred hot spot geometry changed significantly leading to very different posterior mass and radius values, $M \sim 1.7 M_{\odot}$ and $R \sim 14.5$ km, compared to the previously reported results in Refs. [40, 42]. However, the authors emphasized that their results are not yet robust due to computational limitations and, as explained in Ref. [103], the inclusion of XMM-Newton data could introduce systematic biases. In contrast, a reanalysis of PSR J0740+6620 based on more NICER data [280, 281] found consistent results with the previous findings in Refs. [41, 43].

Very recently, preliminary results have been reported from an analysis of the latest NICER measurements of the pulsar PSR J0437-4715 [46]. This pulsar, the brightest object observed by the NICER telescope so far, is located near an even more luminous active galactic nucleus. The object forms a binary system with a white dwarf companion, facilitating the use of an accurate mass prior based on Shapiro time delay measurements [282]. The following preliminary masses and radii were reported at the 68% level:

$$\begin{aligned} \text{PSR J0437-4715} \quad R &= 11.36_{-0.63}^{+0.95} \text{ km} , \\ M &= 1.418 \pm 0.037 M_{\odot} . \end{aligned} \quad (3.72)$$

The inferred mass of PSR J0437-4715 is comparable to that of PSR J0030+0451 as reported in Refs. [40, 42], although the newly measured radius is quite smaller. Nevertheless, both measurements are consistent within their 68% credible intervals⁵. The posterior samples from this analysis are not yet publicly available, therefore we model the two-dimensional mass-radius distribution using a product of two skewed Gaussian distributions:

$$p(M, R|\mathcal{D}) \propto \mathcal{N}_{\text{skew}}(M, a_M, \mu_M, \sigma_M) \mathcal{N}_{\text{skew}}(R, a_R, \mu_R, \sigma_R) , \quad (3.73)$$

where the distribution $\mathcal{N}_{\text{skew}}(x, a, \mu, \sigma)$ corresponds to a combination of a standard Gaussian with the conventional error function:

$$\mathcal{N}_{\text{skew}}(x, a, \mu, \sigma) = \mathcal{N}(x, \mu, \sigma) \left(1 + \operatorname{erf} \left[\frac{a}{\sqrt{2}} \frac{x - \mu}{\sigma} \right] \right) . \quad (3.74)$$

The parameters for the mean μ , standard deviation σ , and skewness a are adjusted to match the reported marginal values in Eq. (3.72). This approximation yields credible levels similar to those reported in Ref. [46]. The likelihood can then be computed using Eq. (3.71), replacing the kernel density estimate with Eq. (3.73). With only preliminary results available so far, we will analyze the impact of the mass-radius measurement of PSR J0437-4715 separately from the other data. As we will later see, while this further refines the inference results, the inclusion of the newest NICER measurement still leads to a consistent picture with the other neutron star data.

A number of other objects are set to be measured by the NICER telescope [67, 276, 283]. Even more sources will be measured by the future missions eXTP [284] and STROBE-X [285], which will be sensitive enough to allow pulse profile modeling for accretion-powered pulsars and thermonuclear burst oscillations [276].

QUIESCENT LOW-MASS X-RAY BINARIES

Low-mass X-ray binaries in quiescence are binary star systems composed of a low-mass donor star transferring material onto a neutron star, which, during periods of low accretion activity, emits minimal X-ray radiation [269, 271]. These binaries are identified in globular clusters where distances, ages, and reddening are relatively well-known [286].

Thermal emissions from quiescent low-mass X-ray binaries (qLMXBs) can be used to determine their masses and radii. The luminosity of a source is given by its distance d and the bolometric flux F :

$$L = 4\pi d^2 F . \quad (3.75)$$

At low temperatures the flux is strongly affected by its passage through the interstellar medium due to photoelectric absorption by electrons. The amount of the absorption depends on the column density of hydrogen atoms along the line of sight, N_H . In the quiescent state, assuming the accretion-heated neutron star surface has a uniform temperature, the luminosity follows the Stefan–Boltzmann law:

$$L = 4\pi R^2 \sigma_{\text{SB}} T_{\text{eff}}^4 , \quad (3.76)$$

where σ_{SB} is the Stefan–Boltzmann constant and T_{eff} the effective surface temperature of the source, which can be extracted by spectral analysis of the X-ray emissions using

⁵ Note that a possible, although less preferred, hot spot configuration in the reanalysis of PSR J0030+0451 [279] obtained very similar mass and radius values as those reported for PSR J0437-4715.

neutron star atmosphere models. Combining Eqs. (3.75) and (3.76), the radius can be determined from

$$R = \sqrt{\frac{d^2 F}{\sigma_{\text{SB}} T_{\text{eff}}^4}}. \quad (3.77)$$

This needs to be corrected to account for the gravitational redshift as well as the light-bending of the thermal photons, both of which depend on the compactness and thus the mass of the neutron star.

In Ref. [141] the information from eight qLMXBs in globular clusters were combined to determine the radius of a neutron star with a characteristic mass of $M = 1.4 M_{\odot}$. In their baseline scenario the authors reported a 68% confidence interval of:

$$R(1.4 M_{\odot}) = 11.4 - 13.1 \text{ km} . \quad (3.78)$$

The assumption of a uniform surface temperature is often motivated by the absence of pulsations in the observed spectra. Temperature inhomogeneities like hot spots would lead to periodic variations in the X-ray flux. However, the absence of pulsations does not guarantee a uniform temperature, as similar behavior could arise from favorable geometries, such as either hot spots or the line of sight being aligned with the rotation axis [160, 287]. Moreover, the assumed atmospheric composition can introduce systematic uncertainty in the thermal X-ray spectra analysis. Hydrogen and helium atmospheres generally provide equally good fits but can yield vastly different radii [269, 271]. Some information about the atmospheric composition can, however, be inferred from the donor star, as the previously accreted matter determines the neutron star's atmospheric ingredients [103, 141].

THERMONUCLEAR BURSTERS

In certain binary star systems containing a neutron star, observations show recurrent bursts of X-ray radiation. These bursts are believed to be triggered by runaway thermonuclear fusion on the accreted envelope of the neutron star due to transfer of hydrogen-rich material from its companion [288]. During the most energetic of these thermonuclear X-ray bursts, the luminosity of the neutron star can exceed the Eddington limit, which represents the theoretical maximum luminosity a star can sustain before radiation pressure overcomes gravity:

$$L_{\text{Edd}} = \frac{4\pi G_N M}{\kappa}. \quad (3.79)$$

Here, κ denotes the radiative opacity that is primarily caused by Thomson scattering. When the Eddington limit is surpassed, the photosphere, i.e., the visible emitting surface of the star, becomes unbound and undergoes a rapid expansion. The photospheric radius expansion leads to a decrease in temperature until the photosphere collapses back onto the neutron star's surface. The collapse triggers a subsequent heating phase, resulting in a characteristic double-peaked structure in the observed burst spectra [289]. The moment when the photosphere collapses back onto the surface is referred to as the 'touchdown', and it is hypothesized that the luminosity takes the value of the Eddington limit [271, 290]. By analyzing the subsequent cooling phase and comparing it to theoretical models of passively cooling neutron star atmospheres, it is possible to extract information about the neutron star's mass and radius.

Ref. [291] employed this ‘cooling tail method’ to thermonuclear X-ray bursts of the pulsar 4U 1702-429, leading to the following 68% mass and radius credible intervals:

$$R = 12.4 \pm 0.4 \text{ km} , \quad (3.80)$$

$$M = 1.9 \pm 0.3 M_{\odot} . \quad (3.81)$$

Note that this analysis assumed a uniformly emitting and slowly rotating neutron star [291]. Moreover, the composition of the star’s atmosphere can impact the opacity of the star leading to different radius values.

In several previous studies measurements from quiescent low-mass X-ray binaries and thermonuclear bursters were included in the Bayesian inference [2, 63–65, 91]. However, as discussed above, these data involve lots of specific model features (see also the discussions in Refs. [103, 160]). Therefore, we do not include them in our Bayesian analysis.

CENTRAL COMPACT OBJECTS

The observation of a neutron star with an unusually small mass, a central compact object within the supernova remnant HESS J1731-347, was recently reported [45]. Such central compact objects have weak magnetic fields and nearly constant thermal X-ray emission. Hence, if the star’s distance is determined via parallax, their mass and radius can be extracted from their spectral emissions similar to qLMXBs. In their analysis, the authors found a mass of only $M = 0.77^{+0.20}_{-0.17} M_{\odot}$ as well as a small radius, $R = 10.4^{+0.86}_{-0.78} \text{ km}$ [45]. This low mass is remarkable because it is not clear how neutron stars with masses lower than around $1.17 M_{\odot}$ form based on known neutron star evolution mechanisms that involve supernovae [156]. The previously known lightest neutron star is compatible with this low mass constraint [292]. This in combination with the object’s small radius led the authors to speculate that HESS J1731-347 might be a possible strange star. Because of the absence of pulsations in the measured spectra, the authors assumed a uniform surface temperature as well as a carbon atmosphere. In addition, they assumed that the stellar magnetic field has no surface effects. However, other authors have suggested a non-uniform emission for similar kinds of central compact objects [293]. In that case larger masses and radii might be possible for HESS J1731-347. Despite such obvious model dependence in the interpretation of the data, it is instructive to add HESS J1731-347 as a separate item in our inference analysis, just in order to explore what its impact would be on the overall picture. The mass-radius posterior from Ref. [45] is again available as samples, so we can compute the likelihood in a way similar to the NICER analyses in Eq. (3.71).

3.4.3 TIDAL DEFORMABILITY MEASUREMENTS

GRAVITATIONAL WAVES

The acceleration of massive objects results in distortions of spacetime known as *gravitational waves* which propagate at the speed of light. Coalescing binary neutron stars are a potent source of gravitational waves, due to the strong acceleration of their masses during the inspiral. Employing laser interferometry techniques, the ground-based detectors of the LIGO and Virgo Scientific Collaborations are capable of measuring the exceedingly small distance changes caused by gravitational waves

The inspiral and coalescence of two compact objects are well described by general relativity. Theoretical waveforms for binary neutron star merger events can be derived from Einstein’s equation through a post-Newtonian expansion [170, 294]. During the

early stages of the inspiral, the time evolution of the gravitational wave frequency is predominantly determined by the chirp mass, $M_{\text{chirp}} = (M_1 M_2)^{3/5} (M_1 + M_2)^{-1/5}$ [295]. As the orbital separation between the two compact objects decreases, the gravitational wave frequency increases and relativistic effects depending on the mass ratio, M_2/M_1 , become more important. As discussed in Sec. 2.1.3, the intense gravitational fields cause the neutron stars to tidally deform, thereby accelerating the coalescence. At fifth-order in the post-Newtonian expansion, the gravitational waveform depends on a mass-weighted combination of the tidal deformabilities of the two neutron stars:

$$\bar{\Lambda} = \frac{16}{13} \frac{(M_1 + 12M_2)M_1^4 \Lambda_1 + (M_2 + 12M_1)M_2^4 \Lambda_2}{(M_1 + M_2)^5}. \quad (3.82)$$

Hence from gravitational wave measurements a Bayesian posterior for the masses (M_1, M_2) and tidal deformabilities (Λ_1, Λ_2) can be inferred. Higher order terms in the gravitational waveform would allow to extract the tidal deformabilities of each star individually. They are, however, unlikely to be detected given the sensitivity of the current detectors [160]. So far, two possible binary neutron star merger events, GW170817 [38] and GW190425 [39], were detected, yielding the following constraints at the 90% level:

$$\text{GW170817} \quad \bar{\Lambda} = 300_{-230}^{+420}, \quad (3.83)$$

$$\text{GW190425} \quad \bar{\Lambda} \leq 600. \quad (3.84)$$

Notice that different analyses of the gravitational wave data produced slightly changed results [296]. The first one of these events (GW170817) was further evaluated together with electromagnetic signals [70, 297, 298]. The following masses and tidal deformabilities of the individual neutron stars in the binary were extracted in Ref. [298] at the 90% level:

$$\begin{aligned} M_1 &= 1.46_{-0.09}^{+0.13} M_\odot & \Lambda_1 &= 255_{-171}^{+416}, \\ M_2 &= 1.26_{-0.12}^{+0.09} M_\odot & \Lambda_2 &= 661_{-375}^{+858}. \end{aligned} \quad (3.85)$$

Using universal relations in Ref. [295] the tidal deformability of a $1.4 M_\odot$ neutron star was extracted based on the information from the merger event GW170817 (90% level):

$$\Lambda(1.4 M_\odot) = 190_{-120}^{+390}. \quad (3.86)$$

We can again approximate the underlying probability distribution of the available posterior samples via kernel density estimation. For a given set of parameters θ , solving the TOV equations in combination with the equations for the tidal deformability yields the relationship $\Lambda(M, \theta)$. Masses larger than the maximum mass of a given EoS, $M > M_{\text{max}}(\theta)$, correspond to black holes, in which case $\Lambda(M, \theta)$ is set to zero. The likelihood is given by the integral over both neutron star masses, with the tidal deformabilities given by $\Lambda(M, \theta)$:

$$\text{Pr}(\mathcal{D}_{\text{Mass-tidal}} | (M, \Lambda)(\theta)) = \int dM_1 \int dM_2 \text{KDE}(M_1, M_2, \Lambda(M_1, \theta), \Lambda(M_2, \theta)). \quad (3.87)$$

Following Ref. [72] we do not use the mass prior here, as we cannot start a priori from assuming the events to be neutron star-neutron star mergers, but should also allow for the principal possibility of neutron star-black hole binaries.

The chirp mass of the GW170817 event inferred in Ref. [38] has a very small uncertainty, $M_{\text{chirp}} = 1.186 \pm 0.001 M_\odot$, hence some analyses assumed it to be fixed. For a

given M_1 this allows to determine M_2 and effectively removes one of the integrations. However, this does not work for GW190425 where the chirp mass has a larger uncertainty, $M_{\text{chirp}} = 1.44 \pm 0.02 M_{\odot}$. Moreover, determining M_2 via M_{chirp} neglects the M_1 dependence of the chirp mass. We perform the double integration in Eq. (3.87) over both M_1 and M_2 without resorting to a fixed chirp mass.

Together with the neutron star merger event GW170817 the short gamma-ray burst GRB170817A and the kilonova AT2017gfo were detected [299]. The gamma-ray burst is believed to originate from the launch of a relativistic jet, whose interaction with the interstellar medium produced a long afterglow [300]. The material ejected during the merger is highly neutron-rich, creating favorable conditions for rapid neutron capture (r-process) nucleosynthesis. The heavy elements formed in this way undergo radioactive decay, heating the ejected material and causing it to emit electromagnetic radiation visible as a kilonova [301]. Some recent Bayesian analyses include information about this kilonova [70, 78, 80, 88], which however introduces a series of model assumptions and consequently raises the systematic uncertainties. For example, in Ref. [103] two different kilonova models led to quite different posterior results. Accordingly, we do not include information from the electromagnetic counterparts of GW170817 in our Bayesian analysis.

The sensitivity of laser interferometric gravitational wave detectors decreases at higher frequencies, which prohibited the extraction of gravitational wave signals from the merger product produced in the GW170817 event. Nevertheless, theoretical models of short gamma-ray burst emission from binary neutron star mergers suggest a subsequent collapse into a black hole [302]. However, the observed properties of the red and blue kilonova ejecta indicate that the collapse did not occur immediately. Instead, it seems likely that the post-merger object could have been a hyper-massive neutron star that was temporarily stabilized against gravitational collapse by differential rotation [302–304]. If this interpretation is valid, it implies a lower upper bound on the maximum mass of non-rotating neutron stars.

Additional merger events will hopefully be detected during the fourth and fifth observation runs of LIGO, Virgo and KAGRA [72, 305]. Furthermore, the next generation of ground-based detectors, i.e., the Einstein telescope [306] and Cosmic Explorer [307], promises a substantial enhancement in the sensitivity of gravitational wave detections. These future facilities are projected to detect multiple binary neutron star mergers daily [308]. In addition, future detectors could enable the observation of gravitational waves from damped oscillations within neutron stars, known as quasi-normal modes (see, e.g., Refs. [309, 310] and references therein), which are highly sensitive to internal discontinuities potentially caused by phase transitions [311]. Moreover, these next-generation instruments are anticipated to detect continuous gravitational waves from rotating neutron stars, facilitating the determination of their moment of inertia [160, 312], a quantity which is sensitive to the neutron star equation of state. Note that the first moment of inertia measurement of a neutron star is expected to become available in the coming years through radio timing observations of the neutron star PSR J0737-3039A in a double pulsar system (see, e.g., Ref. [313] and references therein).

3.4.4 NUCLEAR DATA

Elastic scattering of polarized electrons off heavy nuclei can be used to measure the parity-violating asymmetry, which is sensitive to the interference between electromagnetic interactions via photon exchanges and weak interactions mediated by Z boson exchanges [314]. These interactions differ significantly between protons and neutrons, due to the neutrons' neutral electric charge and the protons' small weak charge. Consequently, the

magnitude of the parity-violating asymmetry depends on the spatial distributions of protons and neutrons within the nuclei, enabling the determination of the difference between the radii of neutron and proton distributions called the neutron skin thickness. This parameter is strongly correlated with the equation of state around nuclear saturation density, $n \sim n_0$, which influences the radius of a typical $1.4 M_\odot$ neutron star [315].

The thick neutron skin thickness of ^{208}Pb extracted in the PREX II measurement [314] suggests a stiff equation of state at small densities [82, 315, 316]. This is in tension with other laboratory probes [315] including the thin neutron skin thickness of ^{48}Ca , as determined by the CREX collaboration [317], which implies a softer EoS in agreement with ChEFT predictions [318]. The tension between these measurements may indicate potential systematic uncertainties in deriving the EoS from the parity-violating asymmetry [103]. Therefore, we do not include these measurements in the total likelihood.

Properties of nuclear matter under extreme density conditions can also be probed through heavy-ion collisions [88, 157, 319]. In these experiments, the collision of ions traveling at relativistic speeds creates a fireball of highly compressed and extremely hot matter that undergoes rapid expansion and emits fast-moving particles. A key observable in these experiments is the elliptic flow, which quantifies the anisotropy of the emitted particles and is particularly sensitive to the early stages of the system's evolution. Notably, at the FOPI [320] and ASY-EOS [321] experiments, gold nuclei were collided at relativistic energies in the range of $0.4 - 1.5$ GeV per nucleon. The compressibility of symmetric nuclear matter could be deduced by comparing the elliptic flow data from the FOPI experiment with predictions from theoretical transport models [320]. These models simulate the non-equilibrium evolution of nuclear matter from the initial impact through the various stages of the collision [157]. In the ASY-EOS experiment, differences in the elliptic flow between neutrons and charged particles (consisting mainly of protons) allowed the determination of the nuclear symmetry energy, which reflects the energy required to convert protons into neutrons in nuclear matter, up to densities about $2n_0$ [321]. These insights from heavy-ion collisions have led to constraints on the pressure of neutron star matter that are comparable to those derived from ChEFT [88, 103]. However, because of the model dependence of the theoretical transport descriptions we do not include information from heavy-ion collisions in our Bayesian analysis.

3.4.5 THEORY CONSTRAINTS

LOW-DENSITY CONSTRAINT FROM CHIRAL EFFECTIVE FIELD THEORY

In many other recent Bayesian studies, the ChEFT constraints introduced in Sec. 2.2.2 were implemented as a prior. However, the ChEFT framework with its low-energy constants and uncertainty measures should be considered as representing a large variety of empirical nuclear physics data, in the same general category as the astrophysical data. There is, in principle, no reason to trust the uncertainty estimates of ChEFT more than those of the astrophysical data. Therefore, we employ the ChEFT information as a likelihood instead of a prior. The likelihood treatment permits a balancing between the constraints from nuclear physics and astrophysics and allows for a rigorous and statistically consistent Bayes factor analysis.

With the Gaussian process used in Refs. [60, 61], the speed of sound displayed in Fig. 2.4 is normally distributed at each density n with mean value $\langle c_s^2(n) \rangle$ and standard deviation $\sigma(n)$. We employ the ChEFT results at discrete densities n_i starting from the BPS crust and extending up to a maximum density for the applicability of the effective field

theory, $n_i \in [0.5 n_0, n_{\text{ChEFT}}]$. The ChEFT likelihood is then computed via Bayesian linear regression, that is the total likelihood is given by the product of the Gaussian likelihoods at each n_i :

$$\begin{aligned} p(\mathcal{D}_{\text{ChEFT}} | c_s^2(n, \theta)) &\propto \prod_i \exp \left[-\frac{1}{2} \left(\frac{\langle c_s^2(n_i) \rangle - c_s^2(n_i, \theta)}{\sigma(n_i)} \right)^2 \right] \\ &= \exp \left[-\frac{1}{2} \sum_i \left(\frac{\langle c_s^2(n_i) \rangle - c_s^2(n_i, \theta)}{\sigma(n_i)} \right)^2 \right], \end{aligned} \quad (3.88)$$

where we assign each density the same prior weight. The speed-of-sound constraint of Refs. [60, 61] is continuous in n . Therefore, we choose to replace the sum by an integral:

$$p(\mathcal{D}_{\text{ChEFT}} | c_s^2(n, \theta)) \propto \exp \left[-\frac{1}{2} \int_{0.5 n_0}^{n_{\text{ChEFT}}} dn \left(\frac{\langle c_s^2(n) \rangle - c_s^2(n, \theta)}{\sigma(n)} \right)^2 \right]. \quad (3.89)$$

We have checked that this likelihood leads to posterior credible bands for the sound velocity that are very similar to those in Refs. [60, 61]. Following the results of the analysis in Ref. [73], we choose a conservative maximum density for the applicability of ChEFT as $n_{\text{ChEFT}} = 1.3 n_0$. At higher densities the N²LO and N³LO results become increasingly different, hinting towards possible convergence issues. We will examine the impact of this choice in the later analysis.

HIGH-DENSITY MATCHING TO PERTURBATIVE QCD

As described in Sec. 2.2.3, at asymptotically high densities, pQCD provides a constraint to be matched in extrapolations far beyond the conditions realized in neutron star cores. Here, we use the likelihood introduced in Refs. [62, 86] to include this matching in our inference analysis. The partial N³LO pQCD results of Ref. [49] depicted in Fig. 2.5 are taken to be valid at the chemical potential $\mu_{\text{pQCD}} = 2.6 \text{ GeV}$, with corresponding density n_{pQCD} and pressure P_{pQCD} . For any point of a given equation of state, n_{NS} , P_{NS} and μ_{NS} , it must be possible to connect to the asymptotic pQCD constraint via a causal and thermodynamically stable interpolation. Because the squared speed of sound (derived in the grand canonical approach from $P(\mu)$) is causally limited,

$$c_s^2 = \left(\frac{\mu}{n} \frac{\partial n}{\partial \mu} \right)^{-1} \leq 1, \quad (3.90)$$

a minimal slope of the function $n(\mu)$ is determined for any specific EoS:

$$\frac{\partial n}{\partial \mu} \geq \frac{n}{\mu}. \quad (3.91)$$

Demanding that it should be possible to connect a point in the neutron star range, $(\mu_{\text{NS}}, n_{\text{NS}}, P_{\text{NS}})(\theta)$, to μ_{pQCD} , n_{pQCD} and P_{pQCD} leads to the integral constraint

$$\int_{\mu_{\text{NS}}}^{\mu_{\text{pQCD}}} d\mu \ n(\mu) = P_{\text{pQCD}} - P_{\text{NS}} = \Delta P. \quad (3.92)$$

It was shown in Ref. [62] that the requirements of causality and thermodynamic stability imply the following minimum and maximum values for ΔP :

$$\Delta P_{\min} = \frac{\mu_{\text{pQCD}}^2 - \mu_{\text{NS}}^2}{2\mu_{\text{NS}}} n_{\text{NS}} , \quad (3.93)$$

$$\Delta P_{\max} = \frac{\mu_{\text{pQCD}}^2 - \mu_{\text{NS}}^2}{2\mu_{\text{pQCD}}} n_{\text{pQCD}} . \quad (3.94)$$

Accordingly, the pQCD likelihood is equal to one if the difference ΔP is within these two values and zero otherwise:

$$p(\mathcal{D}_{\text{pQCD}} | \Delta P(\theta)) = \begin{cases} 1 & \text{if } \Delta P(\theta) \in [\Delta P_{\min}(\theta), \Delta P_{\max}(\theta)] \\ 0 & \text{else} \end{cases} . \quad (3.95)$$

The pQCD results still depend on the renormalization scale X . We follow Refs. [62, 86] and take the logarithmic average over $X \in [1/2, 2]$. Each EoS is constrained by neutron star data only up to the respective maximum central chemical potential $\mu_{c,\max}$, density $n_{c,\max}$ and pressure $P_{c,\max}$. As in Refs. [94, 137] we verify at that point, $(\mu_{\text{NS}}, n_{\text{NS}}, P_{\text{NS}}) = (\mu_{c,\max}, n_{c,\max}, P_{c,\max})$, whether a causal and thermodynamic interpolation to the asymptotic pQCD constraint exists. Note that the authors of Refs. [62, 86] chose $n_{\text{NS}} = 10 n_0$ instead, together with the corresponding chemical potential and pressure, well above the central densities reached in neutron stars. We will analyze the impact of this choice in the later analysis.

TOTAL LIKELIHOOD

Summarizing the preceding sections, the full procedure to obtain posterior credible bands for neutron star properties consists of the following steps: first a set of parameters θ is sampled from the prior $p(\theta)$. We need up to 900,000 samples (with less samples necessary for parametrizations with fewer parameters) in order to generate statistically solid results that cover enough probability mass such that they remain stable after a further increase in the number of samples. For this sampled set of parameters we compute the speed of sound for the respective parametrization and then the EoS, $P(\varepsilon, \theta)$, using Eq. (2.12). Given the equation of state we can numerically solve the coupled system of differential equations introduced in Secs. 2.1.2 and 2.1.3 for $(M, R, \Lambda)(\theta)$. The total likelihood can then be determined as the product of the individual likelihoods for the different measurements and constraints:

$$p(\mathcal{D} | \theta) \propto p(\mathcal{D}_{\text{Mass}} | M(\theta)) p(\mathcal{D}_{\text{Mass-radius}} | (M, R)(\theta)) p(\mathcal{D}_{\text{Mass-tidal}} | (M, \Lambda)(\theta)) \\ \times p(\mathcal{D}_{\text{ChEFT}} | c_s^2(n, \theta)) p(\mathcal{D}_{\text{pQCD}} | \Delta P(\theta)) . \quad (3.96)$$

The prior probability distribution weighted with the above likelihood yields the posterior probability distribution, $p(\theta | \mathcal{D})$, for the parameters. We can then marginalize over this posterior probability distribution to compute the median as well as the highest density credible intervals at the 68% and 95% level for different neutron star properties, as well as the credible bands at different levels as explained in Sec. 3.2.

The set consisting of Shapiro time delay data, NICER measurements and the information from gravitational wave events, together with ChEFT and pQCD constraints, is denoted

‘Previous’ in the following. This serves to study the impact of the three new observations, i.e., the black widow (BW) mass, the supernova remnant (HESS) and the newest NICER (NN) mass-radius data. This nomenclature is also displayed in Table 3.3, where we summarize all data used in the conventional Bayesian analysis.

Data and constraints		
	PSR J1614–2230	$M = 1.908 \pm 0.016 M_{\odot}$ [34]
	PSR J0348+0432	$M = 2.01 \pm 0.04 M_{\odot}$ [35]
	PSR J0030+0451	$M = 1.34^{+0.15}_{-0.16} M_{\odot}$ $R = 12.71^{+1.14}_{-1.19}$ km [40]
Previous	PSR J0740+6620	$M = 2.072^{+0.067}_{-0.066} M_{\odot}$ $R = 12.39^{+1.30}_{-0.98}$ km [41]
	GW170817	$\bar{\Lambda} = 300^{+420}_{-230}$ [38]
	GW190425	$\bar{\Lambda} \leq 600$ [39]
	ChEFT	[60, 61]
	pQCD	[49, 62, 86]
BW	PSR J0952-0607	$M = 2.35 \pm 0.17 M_{\odot}$ [44]
HESS	HESS J1731-347	$M = 0.77^{+0.20}_{-0.17} M_{\odot}$ $R = 10.4^{+0.86}_{-0.78}$ km [45]
NN	PSR J0437-4715	$M = 1.418 \pm 0.037 M_{\odot}$ $R = 11.36^{+0.95}_{-0.63}$ km [46]

TABLE 3.3. Data and constraints used in the conventional Bayesian inference analysis. We examine the impact of the new black widow (BW), supernova remnant (HESS) and NICER (NN) data separately from the previously available data (Previous). All listed results are at the 68% level except for the $\bar{\Lambda}$ results based on the gravitational wave measurements which are at the 90% level.

3.5 SIMULATION-BASED INFERENCE

As explained in Sec. 3.2.1, the likelihoods for the equation of state parameters given the observed detector data $p(\theta|\mathcal{D})$ are not analytically available. Therefore, in our conventional Bayesian analysis described in the previous sections we follow the traditional approach to carry out the inference in two steps, where first exterior neutron star properties are extracted from observed astrophysical detector data, followed by EoS inference from a set of stellar masses, radii and tidal deformabilities [139, 140]. This two-step procedure requires additional assumptions including the assumption that the posteriors derived from the detector observations can be used as likelihoods (see Eq. (3.51)) and that the posterior probability distributions can be represented using kernel density estimates or Gaussians. Recently it has been demonstrated that a single-step inference is possible while fully propagating uncertainties [127, 128]. In their studies, the authors inferred the neutron star equation of state directly based on spectroscopic measurements from the thermal surface emission of low-mass X-ray binaries in quiescence. As explained in Sec. 3.4.2, the emitted stellar spectra depend on the star’s mass and radius, and are also affected by the stellar distance, hydrogen column, and effective surface temperature. Uncertainties on

these *nuisance parameters* must be accounted for when making a robust estimation of the equation of state. While the advancements of Refs. [127, 128] allow access to the posterior of the EoS parameters by marginalizing nuisance parameters, the computational expense involved has hindered access to the full posterior.

In this section, we introduce a novel method that provides access to the full posterior distribution of the equation of state and the nuisance parameters directly from telescope spectra of quiescent low mass X-ray binaries⁶. For that we employ a recently developed simulation-based inference technique known as *neural likelihood estimation* (NLE) [144] in which normalizing flows use samples of simulated X-ray telescope spectra to learn the likelihood of an observation as a function of the EoS and nuisance parameters. The direct inference of the neutron star EoS from telescope spectra via neural likelihood estimation in contrast to the traditional two-step inference is shown schematically in Fig. 3.4.

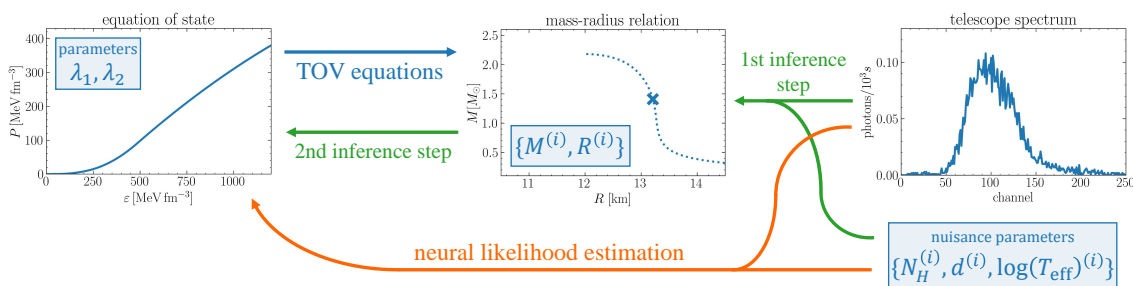


FIG. 3.4. Traditional inference of EoS parameters (left) from telescope spectra (right) is done by first inferring intermediate mass-radius constraints (green arrows), involving additional implicit assumptions. In contrast, neural likelihood estimation allows for inference of the EoS directly from telescope spectra (orange arrow), robustly accounting for uncertainties. Figure taken from Ref. [243].

3.5.1 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. [127, 128], generated using a spectral parametrization with two parameters $\theta = (\lambda_1, \lambda_2)$ according to the prior in Sec. 3.3.3, are used to facilitate direct comparison to previous methods.

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 by fitting the observed spectrum to well-motivated theoretical models, which make assumptions about the source like its atmospheric composition, magnetic field, or temperature [269, 271, 322]. The open-source software XSPEC contains many such models and is widely regarded as the state-of-the-art for spectral fitting [323] and can additionally be used in the generation of simulated spectra. The spectra used in this study were generated using the NSATMOS model in

⁶ As discussed in Sec. 3.4.2, the analysis of qLMXBs involves some model assumptions. We nevertheless focus on thermal X-ray spectra in our analysis to demonstrate the feasibility of the approach and to facilitate a direct comparison with previous works [127, 128]. As discussed later, simulation-based inference methods can, in principle, be similarly applied to other astrophysical data with less systematics, e.g., from NICER or gravitational waves. This section closely follows Ref. [243].

XSPEC, a hydrogen atmospheric model that includes effects of thermal conduction by electrons and self-irradiation due to the gravitational bending of photons [322]. Beyond the 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 [65, 322] and to Poisson noise corresponding to an observation time of 100 ks.

NUISANCE PARAMETERS

Each simulated X-ray spectrum from the NSATMOS model depends on the stellar mass, radius, and three additional parameters: the effective temperature of the surface, T_{eff} , the distance to the star, d , and the hydrogen column, N_H , which parametrizes the reddening of the spectrum by the interstellar medium. These additional neutron star properties are deemed nuisance parameters as they are not directly related to the properties of matter inside the cores of neutron stars but can still greatly impact the observed (and simulated) spectra (see Sec. 3.4.2 for more details on how to extract the masses and radii of qLMXBs). Values for the nuisance parameters are sampled from ranges with distributions motivated by observation as described in Refs. [141, 324]; details of each range are given in Table 3.4.

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_{eff}	exponential	[1, 2] 10^6 K ¹

¹This corresponds to a uniform distribution of $\log(T_{\text{eff}})$ in the range [6, 6.3].

TABLE 3.4. Distributions and ranges of the equation of state parameters (λ_1 and λ_2) and nuisance parameters (N_H , d , and $\log T_{\text{eff}}$) used to generate the samples of simulated neutron stars. See text for parameter definitions.

The values of nuisance parameters can be informed by independent observations, which can vary significantly from star to star and provide prior information. For example, by measuring the redshift of nearby stars one can get an estimate for the distance d of a neutron star to the telescope. We consider three scenarios for nuisance parameter priors, referred to as ‘true’, ‘tight’, and ‘loose’ [127, 128]. 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 3.5 for a description of the prior widths for the three scenarios.

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

TABLE 3.5. Prior distributions on nuisance parameters under three scenarios, ‘true’, ‘tight’, and ‘loose’. Shown are the widths of the Gaussian priors.

3.5.2 PREVIOUS WORK

As discussed in Sec. 3.2, the likelihood, $p(s|\theta)$, of the spectra s given the EoS parameters θ is not analytically known. Traditional methods that infer the EoS from telescope spectra approximate the unknown likelihood in a two-step method [139, 140], 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 [127], occasionally multi-modal [141] contours in the posterior probabilities of stellar mass and radius. Subsequently, these posterior distributions, approximated with a kernel density estimator, are used as likelihoods in a second step to infer the EoS parameters similar to Eq. (3.71). This is valid only if the (M, R) priors used in the inference are sufficiently flat [70, 139].

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. [108–110], 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. The neural network architecture also assumes a fixed number of neutron stars and is not guaranteed to be permutation invariant. Consequently, when new measurements become available, the neural network will need to be retrained. Several other machine learning methods follow a similar approach with varying architectures [111–126]. Thus, they are subject to the same simplifying assumptions. The mass and radius standard deviations can be estimated directly from XSPEC by varying the nuisance parameters [127]; 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 the EoS parameters directly from a set of stellar spectra [127] 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. [128] 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_{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, and they are summarized below:

1. **NN(M, R via XSPEC)**: XSPEC itself is used to predict the M and R of neutron stars. It is run repeatedly for the same star with random nuisance parameters sampled from the prior, to obtain the standard deviations in the predictions of M and R , ΔM and ΔR . A feed-forward neural network is then trained on these four quantities for ten stars to regress the EoS parameters.
2. **NN(Spectra)**: An uncertainty-aware and permutation-invariant network is trained to regress EoS parameters directly from the X-ray spectrum and assumed nuisance parameters. These nuisance parameters are sampled randomly from their respective priors to quantify the uncertainty of the regressed EoS parameters.

3. **ML-Likelihood_{EoS}**: Intractable elements in the relationship between the neutron star’s EoS to the observed spectra are replaced with two neural networks to compute the EoS posterior, marginalizing over nuisance parameters.

3.5.3 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 (θ) and the nuisance parameters (ν) for a star, $p(s|\theta, \nu)$. This extends the strategy in Ref. [128], 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, $p(\theta, \nu|s)$.

NEURAL LIKELIHOOD ESTIMATION

Neural likelihood estimation (NLE) [144] is a type of simulation-based inference [325, 326] technique, successfully used for inference with gravitational waves [327–329], particle physics [330–332] and cosmology [333–337] 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:

$$q_\Phi(s|\theta, \nu) \approx p(s|\theta, \nu) . \quad (3.97)$$

Normalizing flows, as introduced in Sec. 2.3.3, are used as the density estimator. Once trained, NFs can easily generate new samples as well as estimate the likelihood of a given sample. Specifically, a Masked Autoregressive Flow (MAF) [248] is used for the density estimator q_Φ . MAFs are well suited for our purposes because they are very efficient in computing the probability density, but less efficient in generating samples.

The Kullback–Leibler (KL) divergence, D_{KL} , is a measure of the statistical distance between two probability distributions, i.e., it becomes zero if the distributions are identical. Consequently, to approximate a likelihood distribution, $p(s|\theta, \nu)$, we can fit the parameters Φ of a normalizing flow, $q_\Phi(s|\theta, \nu)$, to minimize the KL divergence between the likelihood and the approximate distribution

$$\begin{aligned} & \arg \min_{\Phi} D_{\text{KL}}(p(s|\theta, \nu) \parallel q_\Phi(s|\theta, \nu)) \\ &= \arg \min_{\Phi} \int ds p(s|\theta, \nu) [\log p(s|\theta, \nu) - \log q_\Phi(s|\theta, \nu)] \\ &\approx \arg \min_{\Phi} \sum_{s_i \sim p(s|\theta_i, \nu_i)} \log p(s_i|\theta_i, \nu_i) - \log q_\Phi(s_i|\theta_i, \nu_i) \\ &= \arg \min_{\Phi} \sum_{s_i \sim p(s|\theta_i, \nu_i)} -\log q_\Phi(s_i|\theta_i, \nu_i) \\ &= \arg \max_{\Phi} \sum_{s_i \sim p(s|\theta_i, \nu_i)} \log q_\Phi(s_i|\theta_i, \nu_i) , \end{aligned} \quad (3.98)$$

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. The key realization

of simulation-based inference is that, while the likelihood $p(s|\theta, \nu)$ may not be tractable, it is implicitly defined by both the TOV equations and XSPEC. The simulated spectra s_i for given EoS and nuisance parameters, (θ_i, ν_i) , are sampled from this exact likelihood distribution:

$$s_i \sim p(s|\theta_i, \nu_i) . \quad (3.99)$$

Thus, according to Eq. (3.98), 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. With the training samples s_i generated using the prescription in Sec. 3.5.1, the `sbi` package [338] 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 [339] over the $N = 5$ best-performing density estimators such that the log-likelihood of any given telescope spectrum s_0 is

$$\log p(s_0|\theta, \nu) \approx \frac{1}{N} \sum_j \log q_{\Phi_j}(s_0|\theta, \nu) . \quad (3.100)$$

POSTERIOR SAMPLING WITH HAMILTONIAN MONTE CARLO

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

$$p(\theta, \nu|s) \propto p(s|\theta, \nu)p(\theta)p(\nu) . \quad (3.101)$$

The prior distribution on the nuisance parameters, $p(\nu)$, is given in Table 3.5, while the prior on the EoS parameters, $p(\theta)$, 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 3.4.

To draw samples from the posterior distribution, methods like importance sampling, 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) [147, 148] 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 E.

Unlike standard approaches, our methodology allows for the simultaneous inference of EoS parameters θ 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.

SCALING TO MULTIPLE OBSERVATIONS

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

$$p(s_{1\dots J}|\theta, \nu_{1\dots J}) = \prod_j p(s_j|\theta, \nu_j) . \quad (3.102)$$

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

$$p(\theta, \nu_{1\dots J} | s_{1\dots J}) \propto \left(\prod_j p(s_j | \theta, \nu_j) p(\nu_j) \right) p(\theta) . \quad (3.103)$$

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 [340–343]. 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.

RESULTS AND DISCUSSION

In this chapter, we present the results obtained from the inference procedures introduced in the previous chapter. We begin with a detailed analysis of the posterior results from the conventional Bayesian inference based on the currently available astrophysical data. We then discuss the implications of these empirical results for our understanding of the structure and composition of dense matter. Finally, we present the outcomes of our novel simulation-based inference method, where the EoS is inferred directly from (simulated) telescope spectra.

4.1 CONVENTIONAL BAYESIAN INFERENCE

This section starts with a presentation of the posterior credible bands (for the sound speed, the EoS, the mass-radius relation and the tidal deformability) as well as the inferred properties of typical neutron stars¹. This is done first based on the ‘Previous’ database specified in Table 3.3. A comparison of the inference results for the different parametrizations introduced in the previous chapter reveals the extent of the prior dependence on the functional form of the chosen EoS parametrization. We then further incorporate the ‘new’ black widow PSR J0952-0607 mass information. Thereafter, we focus on detailed Bayes factor investigations, with special emphasis on low sound speeds in combination with extended phase coexistence regions in neutron star matter. Further issues are the assessment of a twin star scenario and the compatibility of a purely hadronic description of neutron star matter with the data. Next, we discuss the trace anomaly measure with the related question for a possible onset of condensation in dense matter. The impacts of the low-density (ChEFT) and high-density (pQCD) constraints on the inference procedure are carefully examined. The section ends with comments on possible effects on the overall systematics induced by the ultralight HESS J1731-347 supernova remnant and the newest NICER measurement of PSR J0437-4715.

4.1.1 PARAMETRIZATION DEPENDENCE

GAUSSIAN VERSUS SEGMENTS PARAMETRIZATION

By combining the likelihoods for the Shapiro time delay, NICER and gravitational wave data with the constraints from ChEFT and pQCD introduced in Sec. 3.4, we compute the total likelihood (3.96) and, using Eq. (3.48), the posterior probability distribution for the parameters θ .

¹ Parts of the text in this section have been adapted from Refs. [96, 97, 163, 344].

The resulting marginal posterior credible bands for the squared speed of sound, $c_s^2(\varepsilon)$, and for the pressure $P(\varepsilon)$ are displayed in Fig. 4.1 for both the skewed Gaussian as well as the ‘basic’ version of the Segments parametrization with fixed asymptotic behavior introduced in Sec. 3.1. Compared to the prior credible bands in Fig. 3.1, we can see how the posterior bands have become much narrower because of the constraints implied by the observational data. The inferred speed of sound shows the same qualitative behavior for both parametrizations, although with tighter credible bands for the G version. Due to the ChEFT likelihood of Eq. (3.89), the sound velocity remains small at low energy densities, $\varepsilon \lesssim 200 \text{ MeV fm}^{-3}$. Given that the low-density behavior is constrained up to $n_{\text{ChEFT}} = 1.3 n_0$, unlike some previous setups that used $n_{\text{ChEFT}} = 2 n_0$ [60, 96], a steep increase of c_s^2 is possible at densities around twice n_0 , i.e., from energy densities $\varepsilon \sim 250 - 300 \text{ MeV fm}^{-3}$ onward. According to Fig. 4.1, the credible bands of the Segments parametrization extend to smaller sound speeds at low ε and then rise to higher speeds of sound, as this parametrization allows for steeper slopes.

In both parametrizations, the median of $c_s^2(\varepsilon)$ exceeds the conformal limit, $c_s^2 = 1/3$, around $\varepsilon \sim 400 - 450 \text{ MeV fm}^{-3}$. We can quantify the evidence for a violation of the conformal bound by computing the Bayes factor $\mathcal{B}_{c_s^2, \text{max} \leq 1/3}^{c_s^2, \text{max} > 1/3}$ which compares equations of state with maximum squared speed of sound larger than $1/3$ to EoSs with maximum squared sound speed below $1/3$. With Bayes factors of well over 10^3 for both parametrizations, there is extreme evidence that c_s^2 exceeds the conformal bound inside neutron stars. This is consistent with other recent studies [71, 72, 74, 86, 132, 202]. As in Refs. [67–70, 86] we implemented the transition to the neutron star crust discontinuously. This is visible in the speed-of-sound credible bands at very low energy densities but of no quantitative significance.

Going to higher energy densities, the 68% credible bands stay above the conformal limit, whereas the 95% bands allow for sound speeds below this bound, but with very low probability. At intermediate energy densities around 650 MeV fm^{-3} , the 68% credible band for the S version starts to form a plateau that extends up to higher energy densities. With the functional form that is assumed a priori, the G version is not able to form an extended plateau. Instead, its 68% credible band continues rising up to higher energy densities $\varepsilon \sim 900 \text{ MeV fm}^{-3}$, before it turns downwards again.

Due to the integrated nature of the pressure (2.12), the differences between the two parametrizations are less prominently visible compared to those in the speed of sound. Nevertheless, the credibility bands of the Gaussian parametrization are again noticeably narrower. At energy densities $\varepsilon \sim 300 - 750 \text{ MeV fm}^{-3}$, the inferred equations of state turn out to be stiffer than the APR EoS [268]², whereas at higher energy densities $P(\varepsilon)$ increases more slowly as compared to APR. It is a known feature that the APR equation of state becomes too stiff and even violates causality at the highest energy densities. The credible bands of the pressure at an energy density of $\varepsilon = 1 \text{ GeV fm}^{-3}$ agree within uncertainties with the softer EoS extrapolated from pQCD in Ref. [86].

The plateau in c_s^2 observed in the Segments parametrization corresponds to an approximately linear rise of the pressure with increasing energy density. In a double logarithmic depiction of $P(\varepsilon)$ the onset of this behavior is reminiscent of the ‘kink’ noted in Ref. [130]. It is then apparent that such a ‘kink’ is not necessarily a signal of a pronounced softening of the EoS but may just reflect the formation of a plateau in the squared sound velocity.

Using the method described in Sec. 3.2.4, we can compute the Bayes factor comparing

² Available via the CompOSE library [345, 346].




FIG. 4.1. Marginal posterior probability distributions at the 95% and 68% level for the Gaussian (left) and ‘basic’ version of the Segments parametrization (right) for the squared speed of sound c_s^2 and pressure P as a function of energy density ε , inferred from the Shapiro time delay, NICER, ChEFT, pQCD and gravitational wave data listed as ‘Previous’ in Table 3.3. At each ε , there exist 95% and 68% posterior credible intervals for $c_s^2(\varepsilon)$ and $P(\varepsilon)$. These intervals are connected to obtain the posterior credible bands. Similarly, the medians of the marginal posterior probability distributions at each ε are connected (solid lines). Grey bars mark the 68% credible intervals of the central energy densities of neutron stars with masses $M = 1.4 M_\odot$ and $2.1 M_\odot$ in each figure. The dashed black line indicates the value of the conformal bound for the speed of sound and the APR EoS [268] for the pressure.

the evidence for the Gaussian and the Segments parametrization, with the result:

$$\mathcal{B}_{\text{Segments}}^{\text{Gaussian}} = 1.70, \quad (4.104)$$

which indicates that neither parametrization is significantly preferred by the data following the classification in Table 3.1.

The posterior credible bands for the mass-radius relation and the tidal deformability are shown in Fig. 4.2. In the credible bands representation there is no natural ending criterion for the bands. Therefore the results for the mass-radius relation are often just cut after the upper 95% interval of the maximum possible mass [41, 74, 83, 84, 86, 88]. We choose to instead limit the $R(M)$ and $\Lambda(M)$ median by the median of M_{max} , and the credible bands by the upper 68% and 95% credible intervals of the maximum mass. As explained in Sec. 3.2.3, we display credible bands instead of two-dimensional credible regions, because the former are independent of the priors on the variables ε and M .

The results for both parametrizations show a similar trend, with the credible bands for the S version again wider than the G version. The Segments parametrization leads to a maximum mass of $M_{\text{max}} = 2.17_{-0.16}^{+0.09} M_\odot$ and a maximum central density of $n_{c,\text{max}} = 6.2_{-0.6}^{+0.9} n_0$ at

the 68% level, while the Gaussian parametrization leads to a slightly smaller maximum supported mass, $2.13_{-0.10}^{+0.08} M_{\odot}$, and larger maximum central density, $6.6_{-0.6}^{+0.8} n_0$. (If not stated otherwise, from here on we always report medians and 68% credible intervals in the text.)




FIG. 4.2. Marginal posterior probability distributions at the 95% and 68% level for the Gaussian (left) and ‘basic’ version of the Segments parametrization (right) for the radius R and tidal deformability Λ as a function of mass M , inferred from the Shapiro time delay, NICER, ChEFT, pQCD and gravitational wave data listed as ‘Previous’ in Table 3.3. At each M , there exist 95% and 68% posterior credible intervals for $R(M)$ and $\Lambda(M)$. These intervals are connected to obtain the posterior credible bands. Similarly, the medians of the marginal posterior probability distributions at each M are connected (solid lines). The $R(M)$ and $\Lambda(M)$ median and credible bands are plotted until the median, upper 68% or 95% interval of the maximum mass. The mass-radius relation is compared to the marginalized intervals at the 68% level from the NICER data analyses by Riley *et al.* (black) [40, 41] of PSR J0030+0451 and PSR J0740+6620. In addition the 68% mass-radius credible intervals of the thermonuclear burster 4U 1702-429 [291] are displayed as well as the 68% credible interval of $R(1.4 M_{\odot})$ extracted from quiescent low-mass X-ray binaries [141] (blue), both of which are not included in the Bayesian analysis. $\Lambda(M)$ is compared to the masses and tidal deformabilities inferred in Ref. [298] for the two neutron stars in the merger event GW170817 at the 90% level (black) as well as $\Lambda(1.4 M_{\odot})$ at the 90% level extracted from GW170817 [295] (blue).

The posterior mass-radius credible bands are in good agreement with the marginalized 68% credible intervals inferred from the NICER measurements of PSR J0030+0451 and PSR J0740+6620. The $R(M)$ credible band is slightly shifted to smaller radii compared to the NICER data, especially in comparison to the heavier pulsar, because the gravitational wave event GW170817 prefers such smaller radii [66]. Here the balancing between different observables and theoretical constraints becomes visible which requires a statistically well-defined analysis in contrast to simple cuts used, e.g., in Refs. [130–132, 136]. Furthermore, the NICER analyses chose to use central credible intervals. If instead, as in the present work, highest density intervals were used, these intervals would reach to smaller radii.

The 68% credible intervals for the radius of a neutron star with mass $M = 1.4 M_\odot$ agrees well with the radius in Eq. (3.78) extracted from quiescent low-mass X-ray binaries in Ref. [141]. Moreover, the 68% mass and radius credible intervals of 4U 1702-429 in Eq. (3.81) extracted from thermonuclear X-ray bursts in Ref. [291] lie well within the 68% mass-radius credible band. Notably both of these observations were not used as input in the Bayesian inference procedure. We mention that in Ref. [87] quiescent low-mass X-ray binaries and sources of thermonuclear bursts were also found to fit into an overall picture that is consistent with the gravitational wave and NICER data. Finally, if we use the NICER data analyses by Miller *et al.* for the inference procedure instead of the ones by Riley *et al.*, we find very similar results, in the same way as Ref. [74]. So we can restrict ourselves to the latter.

The 90% credible intervals for the tidal deformabilities and masses of the two neutron stars in the merger event GW170817 in Eq. (3.85) extracted in Ref. [298] agree well with the posterior credible band of $\Lambda(M)$ displayed in Fig. 4.2. The 90% credible interval for the tidal deformability of a $1.4 M_\odot$ neutron star in Eq. (3.86) extracted based on the information from the merger event GW170817 in Ref. [295] does agree with the posterior credible bands of $\Lambda(M)$ at $M = 1.4 M_\odot$, which lie, however, at slightly larger tidal deformabilities.

Table 4.1 shows medians and credible intervals for selected properties of neutron stars with characteristic masses $M = 1.4 M_\odot$ or $2.1 M_\odot$, including the central density, the energy density and pressure as well as the radius and tidal deformability. Again these numbers demonstrate agreement within uncertainties between the two parametrizations. The inferred properties for relatively light neutron stars are very similar in both parametrizations. However, the differences become more pronounced for heavy neutron stars with a mass of $2.1 M_\odot$. In this case, both parametrizations yield very similar radii, but the values for the central density and pressure are more different. Nevertheless, they are still consistent within the inferred 68% credible intervals.

At the 68% level the inferred radius of a $1.4 M_\odot$ neutron star, $R = 12.4 \pm 0.5$ km for the S version, agrees with the values found in Ref. [70] for a piecewise polytrope parametrization and a speed of sound model similar to our Gaussian parametrization, while the authors additionally included constraints from modeling of the kilonova AT2017gfo. The 68% credible intervals of the radius and tidal deformability of a $1.4 M_\odot$ neutron star listed in Table 4.1 agree within uncertainties with the results in Ref. [82] which include a theory prediction and the PREX II measurement of the ^{208}Pb neutron skin thickness. Our result for the $1.4 M_\odot$ radius also agrees with the value found in Ref. [88], where the authors additionally incorporated constraints on the EoS deduced from relativistic heavy-ion collisions.

With the ‘Previous’ data set and the Segments parametrization we find $n_c = 4.3_{-1.1}^{+0.8} n_0$ for the central density of a $2.1 M_\odot$ neutron star, comparable to the value for a $2.0 M_\odot$ neutron star reported in Ref. [74]. For the G version we find again a slightly larger central density, $n_c = 4.7 \pm 1.0 n_0$. The radius of a $2.1 M_\odot$ neutron star, $R = 11.7 \pm 0.6$ km in the S version, and very similar in the G version, 11.6 ± 0.6 km, agrees with the value for $R(M = 2.0 M_\odot)$ reported in Ref. [132]. In the Bayesian analysis of Ref. [43], no ChEFT constraint was included at low densities. Their prediction for the radius of a neutron star with mass $M = 1.4 M_\odot$, based on multiple different parametrizations, agrees nonetheless with our result at the 68% level. Their result for the radius of the $2.08 M_\odot$ neutron star is larger compared to our result for the radius of a generic $2.1 M_\odot$ neutron star. However, within the 68% credible intervals the two results are still consistent and the differences can be accounted for by the ChEFT constraint.

		Gaussian		Segments	
		95%	68 %	95%	68%
$1.4 M_{\odot}$	n_c/n_0	2.8 ± 0.6	± 0.3	2.8 ± 0.7	± 0.3
	ε_c [MeV fm $^{-3}$]	461^{+106}_{-102}	$^{+48}_{-51}$	461^{+124}_{-114}	$^{+46}_{-59}$
	P_c [MeV fm $^{-3}$]	63^{+23}_{-17}	$^{+9}_{-10}$	63^{+30}_{-19}	$^{+10}_{-13}$
	R [km]	$12.5^{+0.7}_{-0.8}$	$^{+0.4}_{-0.3}$	$12.4^{+0.7}_{-1.1}$	± 0.5
	Λ	439^{+180}_{-178}	$^{+94}_{-90}$	427^{+181}_{-212}	$^{+110}_{-113}$
$2.1 M_{\odot}$	n_c/n_0	$4.7^{+1.6}_{-1.9}$	± 1.0	$4.3^{+1.9}_{-1.6}$	$^{+0.8}_{-1.1}$
	ε_c [MeV fm $^{-3}$]	865^{+398}_{-409}	$^{+178}_{-278}$	769^{+464}_{-335}	$^{+189}_{-228}$
	P_c [MeV fm $^{-3}$]	299^{+254}_{-201}	$^{+79}_{-164}$	254^{+276}_{-153}	$^{+87}_{-117}$
	R [km]	$11.6^{+1.3}_{-1.1}$	± 0.6	$11.7^{+1.3}_{-1.2}$	± 0.6
	Λ	15^{+26}_{-10}	$^{+5}_{-9}$	18^{+28}_{-12}	$^{+6}_{-11}$

TABLE 4.1. Median, 95% and 68% credible intervals for selected neutron star properties for the Gaussian and ‘basic’ version of the Segments parametrizations given the data set ‘Previous’ from Table 3.3. These are computed from the one-dimensional posterior probability distributions marginalized over all other parameters. Listed are the central density, energy density, pressure, radius and tidal deformability of neutron stars with masses $M = 1.4 M_{\odot}$ and $2.1 M_{\odot}$.

At the current state of investigations with a limited neutron star database and correspondingly large uncertainties, it is still justified to use parametric functional forms as long as they are sufficiently general. In Ref. [89] the authors compare different parametrizations and argue that inferred neutron star results depend on the chosen parametrization. However, in their comparison the primary differences in the inferred equations of state occur at small densities, mainly because of different implementations of the neutron star crust, and in the high density regime not constrained by data. In the intermediate region, $n \sim 1.5 - 6 n_0$, the different parametrizations agree within their uncertainties. One exception is a Gaussian parametrization which, unlike our G version, does not allow skewed Gaussians and is therefore not sufficiently general to reproduce the current astrophysical data, a feature that is already visible from its prior. Our point regarding the stability of inference results with respect to different parametrizations is further supported by the work of Ref. [88] where very similar neutron star properties are found for two qualitatively different parametrizations, namely a Segments parametrization similar to our S version and a piecewise polytrope representation. In the future many more neutron star observations are expected (as explained in Sec. 3.4), which will lead to more precise constraints on the speed of sound. Then, a non-parametric description of the EoS in terms of a Gaussian process [71, 72, 74, 89] or neural network [79, 92] might become preferable.

CONFORMAL LIMIT REACHED FROM ABOVE

A recent analysis based on hard dense loop resummation techniques found that, in contrast to standard perturbative QCD results, the speed of sound reaches the conformal limit from above at asymptotically high baryon densities [207], similar to the situation at large isospin densities [205, 347]. To analyze the impact of this assumed alternative asymptotic

behavior, we repeat our analysis with the same steps as before, but now modified such that the squared speed of sound reaches the conformal limit, $c_s^2 \rightarrow 1/3$, from above and without including the pQCD likelihood introduced in Sec. 3.4.5 to compute the posterior. We restrict ourselves to the Segments parametrization since the descriptive power of the Gaussian parametrization is severely hindered with the changed asymptotic behavior. The resulting posterior credible bands are displayed in Fig. 4.3. Compared to the case with $c_s^2 \rightarrow 1/3$ reached from below in Figs. 4.1 and 4.2, the sound velocities remain relatively unchanged up to energy densities $\varepsilon \lesssim 500 \text{ MeV fm}^{-3}$, implying that both asymptotic behaviors lead to $M = 1.4 M_\odot$ neutron stars with almost exactly the same properties as in Table 4.1. Even at higher energy densities the 68% credible bands look similar. However, when the conformal limit is reached from above, the lower limit of the 95% credible band lies at higher sound speeds. The similarity in the speed of sound translates into the credible bands for $P(\varepsilon)$, $R(M)$ and $\Lambda(M)$. However, the 68% credible $R(M)$ band extends to slightly larger masses. For a $2.1 M_\odot$ neutron star, when the conformal limit is reached from above, we find a central density $n_c = 4.3_{-1.1}^{+1.0} n_0$, a central pressure $P_c = 260_{-125}^{+96} \text{ MeV fm}^{-3}$, a central energy density $\varepsilon_c = 765_{-252}^{+192} \text{ MeV fm}^{-3}$, a radius $R = 11.6_{-0.5}^{+0.7} \text{ km}$ and a tidal deformability $\Lambda = 17_{-11}^{+7}$, again very similar to the previous results in Table 4.2. This means that the description of neutron stars at all mass ranges is to large extent independent of the high density asymptotic behavior as long as the speed of sound is causally connected to the conformal limit.

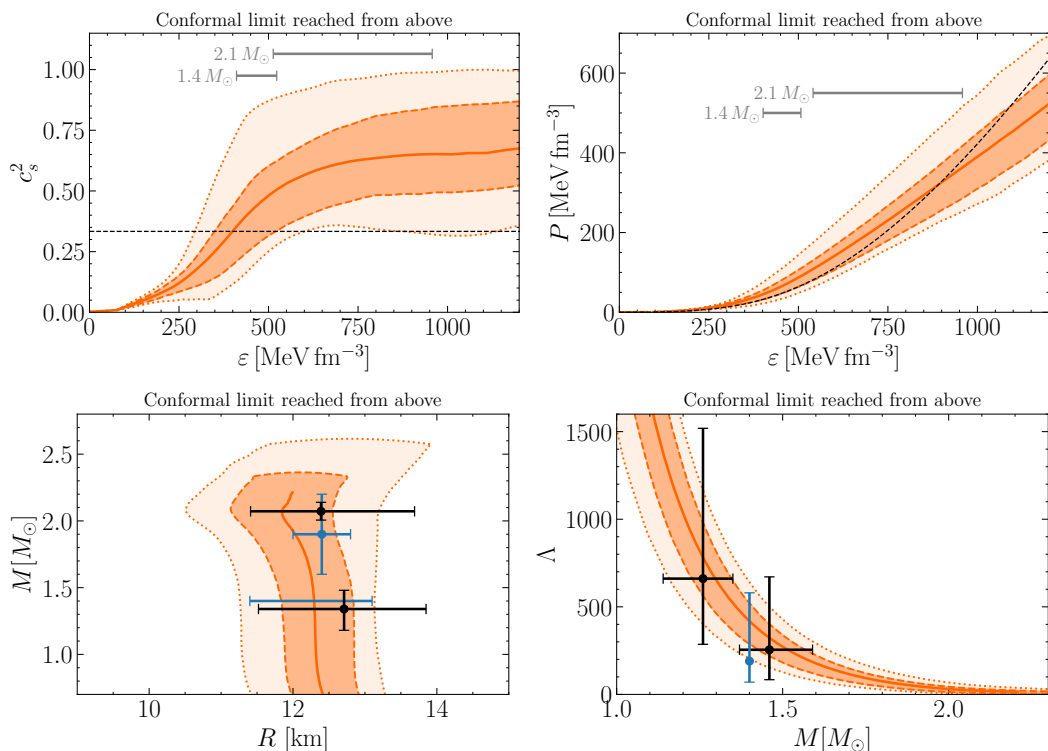


FIG. 4.3. Similar to Figs. 4.1 and 4.2: posterior 95% and 68% credible bands and median for the Segments parametrization with the conformal limit reached asymptotically from above and without including the pQCD likelihood. Shown are the squared speed of sound, c_s^2 and the pressure P as a function of energy density ε , as well as the mass-radius relation $R(M)$ and the tidal deformability $\Lambda(M)$. The credible bands are compared to the same data as in Fig. 4.2.

IMPROVED PARAMETRIZATION

In the previous sections we saw that posterior neutron star properties inferred with the Gaussian and the ‘basic’ version of the Segments parametrization agree within the uncertainties associated with the still limited amount and accuracy of the available astrophysical data. However, the Segments parametrization turns out to be preferred as it leads to slightly larger posterior credible bands, indicative of a less restrictive functional form. Moreover, changing the assumed asymptotic behavior of the sound speed has only a minor impact on the properties of neutron stars. Therefore, in the following we continue with the improved S’ version of the Segments parametrization with two more free parameters and the more general asymptotic behavior introduced in Sec. 3.1. Asymptotically the speed of sound still reaches the conformal limit through the pQCD likelihood introduced in Sec. 3.4.5.

The posterior credible bands inferred from the ‘Previous’ data in Table 3.3, consisting of the Shapiro time delay, NICER, ChEFT, pQCD and gravitational wave data, for the S’ version are displayed in Fig. 4.4. There is a close agreement to the previously employed ‘basic’ version of the Segments parametrization with some minor changes resulting from the altered asymptotic behavior and enhanced functional flexibility. After the ChEFT constraint at small energy densities, the S’ version’s 95% posterior credible bands for the speed of sound allow both for steeper slopes and smaller values than those displayed in Fig. 4.1. Stiffer sound velocities lead to larger radii, as can be seen compared to the S version’s mass-radius credible bands in Fig. 4.2. A comparison of Bayes factors indicates similar evidence for the improved Segments parametrization compared to the S and G versions used in the previous sections.

The minor modifications to the speed of sound also affect the median and credible intervals of the inferred properties of neutron stars with masses of $M = 1.4 M_{\odot}$ and $2.1 M_{\odot}$ collected in Table 4.2. However, these changes are negligible in comparison to the inferred uncertainties, indicating that the inference results are stable with respect to an increase in the number of free segments and to changes in the asymptotic behavior of the parametrization (as already concluded in the previous section). Marginalizing the posterior based on the ‘Previous’ data set with respect to the maximum mass M_{\max} , we find median and 68% credible interval values of $M_{\max} = 2.20^{+0.10}_{-0.16} M_{\odot}$ for the improved Segments parametrization. The corresponding probability distribution is displayed in Fig. 4.5. The 68% credible interval of M_{\max} has slightly increased compared to the previous parametrizations because the S’ version allows for steeper slopes in the speed sound, see Fig. 4.4. Recent studies find similarly large maximum masses [74, 86]. The maximum central neutron star density turns out to be slightly smaller with $n_{c,\max} = 6.0^{+0.7}_{-0.8} n_0$.

4.1.2 IMPACT OF PSR J0952-0607

Including the new mass information of the black widow pulsar PSR J0952-0607 in the database, the maximum non-rotating neutron star mass increases significantly to $M_{\max} = 2.31^{+0.14}_{-0.17} M_{\odot}$, as displayed in Fig. 4.5. The median and the 68% credible interval extend to masses lower than that of the PSR J0952-0607 mass, $2.35 \pm 0.17 M_{\odot}$. This is due to the large mass uncertainty and the rotation correction that we have applied in the analysis. To support such higher masses the sound speed needs to increase more rapidly, implying a stiffer EoS which in turn leads to smaller central densities in the neutron star core. Accordingly, the maximum central density decreases to $n_{c,\max} = 5.6 \pm 0.7 n_0$ with inclusion of the new heavy-mass data.

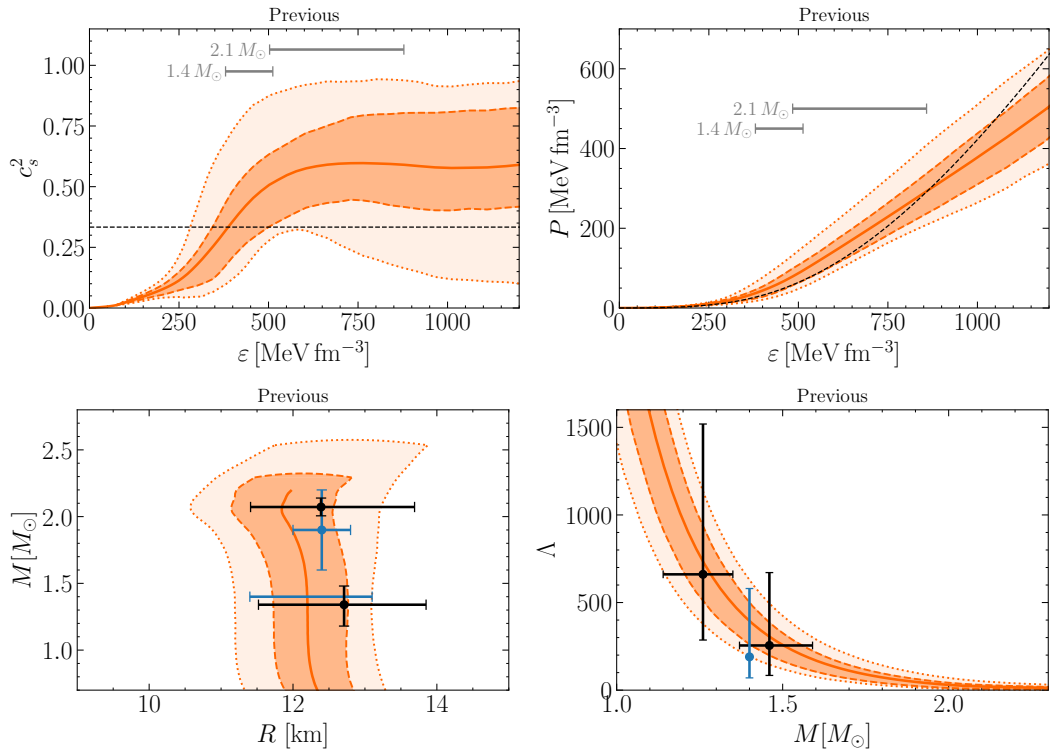


FIG. 4.4. Similar to Figs. 4.1 and 4.2: posterior credible bands are displayed for the squared speed of sound, c_s^2 , and the pressure, P , as a function of energy density ϵ , and the mass-radius relation $R(M)$ and tidal deformability $\Lambda(M)$, inferred from the ‘Previous’ data in Table 3.3, but now using the improved Segments parametrization with two more free parameters and a more general asymptotic behavior. Figure (partly) taken from Ref. [97].




FIG. 4.5. Marginal posterior probability distributions of the maximum neutron star mass M_{\max} . The nomenclature ‘Previous’ refers to the Shapiro time delay, NICER, ChEFT, pQCD and gravitational wave data listed in Table 3.3. The shift induced by adding the (non-rotating) mass information from the black widow pulsar PSR J0952-0607 (BW) is also shown. The prior distribution is nearly uniform over a wide mass range. Figure taken from Ref. [97].

The inclusion of the new information from PSR J0952-0607 has only a marginal impact on the properties of neutron stars with relatively small masses, see Table 4.2. However, for a $2.1 M_\odot$ neutron star, the inclusion of the heavy mass information reduces the central density significantly to $n_c = 3.6 \pm 0.7 n_0$. Similarly, the central energy density ϵ_c and

		Previous		Previous + BW	
		95%	68 %	95%	68 %
1.4 M_{\odot}	n_c/n_0	$2.8^{+0.8}_{-0.7}$	± 0.4	2.6 ± 0.7	$^{+0.3}_{-0.4}$
	ε_c [MeV fm $^{-3}$]	451^{+133}_{-123}	$^{+62}_{-71}$	423^{+118}_{-116}	$^{+56}_{-67}$
	P_c [MeV fm $^{-3}$]	64^{+30}_{-23}	$^{+12}_{-16}$	60^{+28}_{-20}	$^{+11}_{-14}$
	R [km]	$12.2^{+0.9}_{-1.0}$	± 0.5	$12.3^{+0.8}_{-1.0}$	± 0.5
	Λ	396^{+226}_{-197}	$^{+107}_{-127}$	421^{+236}_{-200}	$^{+114}_{-124}$
2.1 M_{\odot}	n_c/n_0	$4.1^{+1.9}_{-1.5}$	$^{+0.8}_{-0.9}$	$3.6^{+1.6}_{-1.3}$	± 0.7
	ε_c [MeV fm $^{-3}$]	716^{+416}_{-326}	$^{+162}_{-213}$	628^{+357}_{-251}	$^{+149}_{-146}$
	P_c [MeV fm $^{-3}$]	225^{+239}_{-134}	$^{+62}_{-110}$	186^{+184}_{-104}	$^{+52}_{-80}$
	R [km]	11.9 ± 1.3	± 0.7	$12.1^{+1.3}_{-1.2}$	$^{+0.6}_{-0.8}$
	Λ	21^{+30}_{-15}	$^{+9}_{-13}$	26^{+30}_{-20}	$^{+10}_{-14}$

TABLE 4.2. Similar to Table 4.1: median, 95% and 68% credible intervals for selected neutron star properties for the improved Segments parametrization given the previously available data and the new information from the black widow (BW) pulsar.

central pressure P_c are reduced, while the radius R and the tidal deformability Λ are slightly increased. A very similar impact is seen for the G and S versions.

With the heaviest observed pulsar in mind, Table 4.3 displays inferred properties of a neutron star with mass $M = 2.3 M_{\odot}$. For such a heavy-mass object the central density is still only $n_c = 3.8^{+0.7}_{-0.8} n_0$, comparable to that of a $2.1 M_{\odot}$ neutron star in Table 4.2. This result is of some significance because it indicates that the baryon densities in the cores of even the heaviest neutron stars are not expected to reach extreme values. We will discuss the implications of this for the composition of neutron star matter later in Sec. 4.2.

		Previous + BW	
		95%	68 %
2.3 M_{\odot}	n_c/n_0	$3.8^{+1.6}_{-1.3}$	$^{+0.7}_{-0.8}$
	ε_c [MeV fm $^{-3}$]	673^{+363}_{-268}	$^{+140}_{-180}$
	P_c [MeV fm $^{-3}$]	237^{+226}_{-134}	$^{+69}_{-104}$
	R [km]	12.3 ± 1.2	$^{+0.7}_{-0.6}$
	Λ	14^{+17}_{-10}	$^{+4}_{-9}$

TABLE 4.3. Same as Table 4.2, but median and credible intervals for a neutron star with mass $M = 2.3 M_{\odot}$ are displayed given the previously available data plus new information from the black widow (BW) pulsar PSR J0952-0607.

In Fig. 4.6 the posterior credible bands are displayed for an inference with the new information from PSR J0952-0607. At first sight the comparison between Figs. 4.4 and 4.6 appears to reveal only marginal differences. But with a more focused view the condition to reach the high mass of the black widow pulsar implies that the speed of sound increases

more stiffly at low energy densities, $\varepsilon \lesssim 500 \text{ MeV fm}^{-3}$. Accordingly, the conformal bound is exceeded at smaller energy densities, and higher speeds of sound are reached earlier compared to the case without the heavy mass data in Fig. 4.4. With inclusion of the new data larger radii are necessary for heavy neutron stars with $M \sim 2.1 M_\odot$, such that the mass-radius relation in Fig. 4.6 features an almost constant radius $R \sim 12.3 \text{ km}$ over the whole mass range. For practical applications, the median values of the squared sound velocity c_s^2 as a function of energy density, as shown in Figs. 4.4 and 4.6, are listed in Appendix C.

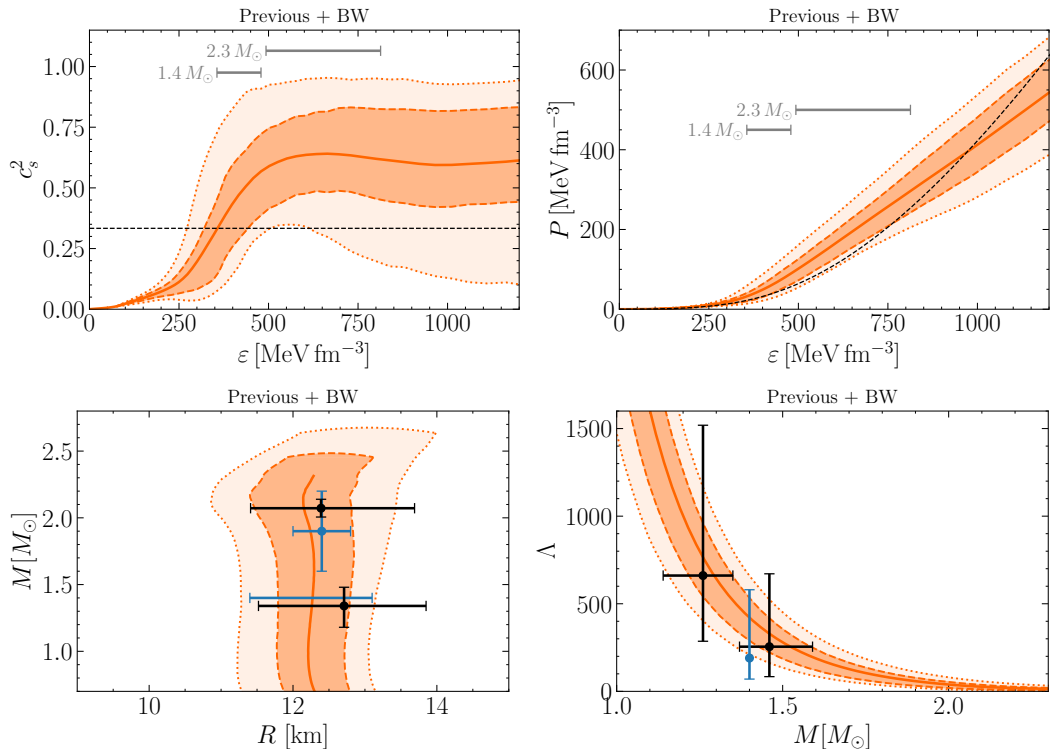


FIG. 4.6. Similar to Fig. 4.4: posterior credible bands are displayed for the improved Segments parametrization for the squared speed of sound, c_s^2 , and the pressure, P , as a function of energy density ε , and the mass-radius relation $R(M)$ and tidal deformability $\Lambda(M)$, but now using the ‘Previous + BW’ data in Table 3.3 including the new PSR J0952-0607 (BW) information. Grey bars mark the 68% credible intervals of the central energy densities of neutron stars with masses $M = 1.4 M_\odot$ and $2.3 M_\odot$, respectively. Figure taken from Ref. [97].

To give an impression of the matter distribution inside neutron stars, the density profiles of objects with masses $M = 1.4 M_\odot$ and $M = 2.1 M_\odot$ are displayed in Fig. 4.7. These profiles are computed using the median of $P(\varepsilon)$ based on the previously available data together with the new information from the black widow pulsar. On the axes of Fig. 4.7 the 68% credible intervals of the central densities and radii of 1.4 and $2.1 M_\odot$ neutron stars listed in Table 4.2 are indicated for comparison. The skewness of the posterior probability distribution makes the central densities and radii in Fig. 4.7 deviate slightly from the median values listed in Table 4.2. As in Fig. 4.6 both neutron stars with masses $M = 1.4 M_\odot$ and $2.1 M_\odot$ have almost equal radii. The density profiles smoothly decrease towards small densities in the outer regions of the stars. In this regime the EoS is governed by the ChEFT constraint.




FIG. 4.7. Density profiles of neutron stars with masses of $M = 1.4 M_{\odot}$ and $M = 2.1 M_{\odot}$. The employed equation of state corresponds to the median of the posterior in Fig. 4.6, i.e., using the previously available and the new data from the black widow pulsar. The bars indicate the 68% credible intervals of the central densities and radii of neutron stars with mass $M = 1.4 M_{\odot}$ (blue) and $2.1 M_{\odot}$ (orange), as listed in Table 4.2. Figure taken from Ref. [97]

THERMODYNAMIC PROPERTIES

In Fig. 4.8 posterior credible bands for the baryon density n as a function of energy density ε based on the previously available data together with the new information from the black widow pulsar are displayed. The energy per particle can be computed as

$$\frac{E}{A} = \frac{\varepsilon}{n} - m_N, \quad (4.105)$$

which is displayed in Fig. 4.8 as a function of baryon density. Here we take $m_N \sim 939.5$ MeV, the neutron mass with minor adjustment for a small proton fraction of $\sim 10\%$ as in the APR EoS.




FIG. 4.8. Posterior 95% and 68% credible bands and medians for the baryon density n , in units of the nuclear saturation density $n_0 = 0.16 \text{ fm}^{-3}$, as a function of energy density ε (left) and the energy per particle $E/A = \varepsilon/n - m_N$ as function of baryon density (right)..

A further quantity of interest is the chemical potential associated with the conserved baryon number,

$$\mu(n) = \frac{\partial \varepsilon(n)}{\partial n} = \frac{P + \varepsilon}{n}, \quad (4.106)$$

where the latter equality is the Gibbs–Duhem relation at zero temperature. The resulting $\mu(n)$ bands in Fig. 4.9 show a steep rise at baryon densities of $n > 2n_0$, indicating that

strongly repulsive forces are at work. The agnostic approach for the inference of the sound speed on which these results are based does not permit distinction between different species of constituents inside neutron stars, so μ represents the total chemical potential from all active degrees of freedom carrying baryon number: $\mu = \sum_i x_i \mu_i$, where $x_i = n_i/n$ is the fraction corresponding to each baryonic species i . However, irrespective of whether these species are nucleons, other baryonic composites, or quarks, the empirically inferred behavior of μ demonstrates that these degrees of freedom must be correlated by strongly repulsive forces as the density increases in order to build up the necessary high pressure in the cores of neutron stars to support two solar masses and beyond. We return to this discussion again at a later stage.




FIG. 4.9. Posterior 95% and 68% credible bands and median for the baryon chemical potential μ as a function of baryon density n in units of the nuclear saturation density, $n_0 = 0.16 \text{ fm}^{-3}$. Figure taken from Ref. [97].

4.1.3 EVIDENCE FOR (OR AGAINST) A STRONG FIRST-ORDER PHASE TRANSITION

A sufficient condition for a first-order phase transition is an equation of state that features a domain of phase coexistence within which a Maxwell construction implies a region of constant pressure. The width of this domain, characterized by $\Delta n/n$ (where n is the density at which the coexistence interval ends), is a measure of the ‘strength’ of the phase transition, i.e., the amount of surface tension between the two coexisting phases. For guidance and comparison, an example of a ‘strong’ first-order transition is the liquid-gas phase transition in symmetric nuclear matter. There, at low temperatures ($T < 15 \text{ MeV}$), the phase coexistence region obtained through a Maxwell construction has a typical width $\Delta n/n \geq 1$ [221, 348, 349].

Starting from a given EoS, $P(\varepsilon)$, the Gibbs–Duhem relation is used to reexpress pressure as a function of density, $P(n)$. The border lines of the 68% and 95% posterior credible bands for $P(n)$ constrain the maximum possible phase coexistence intervals, $(\Delta n/n)_{\text{max}}$, at the corresponding credibility levels. An example is shown in Fig. 4.10. The results in Fig. 4.11 based on the previously available data together with the new information from the black widow pulsar show that these maximum possible coexistence regions are narrow: $(\Delta n/n)_{\text{max}} \simeq 0.2$ (0.3) at the 68% (95%) level³. In fact it turns out that $(\Delta n/n)_{\text{max}}$ is nearly constant as a function of the baryon density n taken at the end point of the possible phase coexistence region. This observation holds throughout the regime relevant

³ For different choices of n , either as the mean density or as the starting point of the phase coexistence interval, $\Delta n/n$ is still $\lesssim 0.2$ within the inferred 68% posterior credible bands.

for neutron stars. We conclude that only weak phase transitions with $\Delta n/n \leq (\Delta n/n)_{\max}$ can still be realized inside neutron stars within the inferred posterior credible bands.

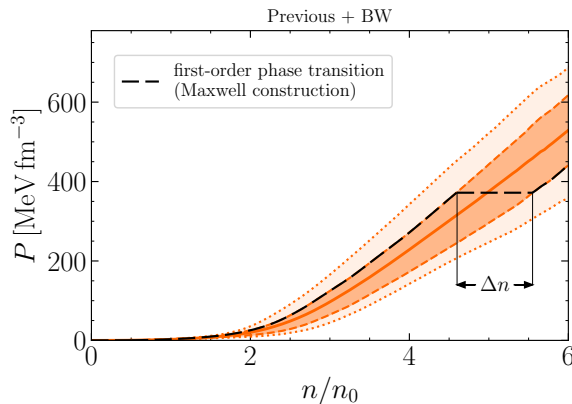


FIG. 4.10. Illustration of the constraint on the maximum width of a Maxwell-constructed coexistence region for a first-order phase transition within the 68% posterior credible band of $P(n)$ with an end density of $n = 5.6 n_0$, in units of the nuclear saturation density, $n_0 = 0.16 \text{ fm}^{-3}$, resulting in a constraint $(\Delta n/n)_{\max} = 0.17$ at this density. The posterior credible bands are derived based on the previously available data together with the new information from the black widow pulsar for the improved version of the Segments parametrization.




FIG. 4.11. Maximum possible phase coexistence interval $(\Delta n/n)_{\max}$ of constant pressure (where n is the density at which the interval ends) extracted from the 68% and 95% posterior credible bands of $P(n)$ in Fig. 4.10 based on the previously available data together with the new information from the black widow pulsar for the improved version of the Segments parametrization. $(\Delta n/n)_{\max}$ is displayed as a function of baryon density n in units of the nuclear saturation density, $n_0 = 0.16 \text{ fm}^{-3}$. Figure taken from Ref. [97].

We note that the study performed in [86] also discussed restrictive constraints on strong first-order phase transitions but did not include the NICER and J0952-0607 black widow pulsar data. Even with this much more limited database, the authors found possible phase coexistence intervals of only $\Delta n/n \lesssim 0.5$ for first-order phase transitions within neutron stars. As can be expected, the additional radius constraints imposed by the NICER data constrain the maximum extension of the phase coexistence region even further. In addition, in Ref. [105], without inclusion of the black-widow pulsar data and different implementations of the ChEFT and pQCD constraints, the author found possible phase

coexistence intervals $\Delta n/n \lesssim 0.4$ for first-order phase transitions that can take place within neutron stars, i.e., do not destabilize the mass-radius sequence.

The detailed behavior of the squared speed of sound is to be seen in a related context. Figs. 4.4 and 4.6 show indications of a shallow maximum, $c_{s,\max}^2$. At the 68% level this maximum takes a value $c_{s,\max}^2 = 0.78_{-0.11}^{+0.18}$ at a baryon density $n(c_{s,\max}^2) = 3.2_{-1.2}^{+0.8} n_0$. The peak in c_s^2 found in Ref. [90] has a similar magnitude and location, although a pronounced peak structure is not seen in our posterior result. In contrast we find that the sound velocity forms a plateau at higher densities. There is no indication of a softening. Still, at the 95% level small sound speeds are not entirely excluded, though the probability of their occurrence is low. Nevertheless, at asymptotically high densities pQCD dictates that the speed of sound reaches the conformal bound $c_s^2 = 1/3$ from below. This implies that at some density beyond the plateau, the speed of sound must turn downward again and reach a minimum, $c_{s,\min}^2$, at some higher density. A fast drop in c_s^2 could potentially indicate the occurrence of a phase transition. The question is whether such a decrease still takes place within the density range of neutron star cores.

To answer this question we specifically perform a Bayes factor analysis to quantify the evidence for a rapid variation of energy density with pressure, corresponding to a low averaged sound speed over the relevant pressure interval. With this aim Bayes factors $\mathcal{B}_{c_{s,\min}^2 \leq 0.1}^{c_{s,\min}^2 > 0.1}$ are computed, comparing the evidence for EoSs with a minimum speed of sound larger than $c_{s,\min}^2 > 0.1$ over EoSs with small $c_{s,\min}^2 \leq 0.1$, the latter possibly indicating a strong first-order phase transition. It is assumed that this minimum is positioned above the maximum located at lower densities, $n(c_{s,\min}^2) > n(c_{s,\max}^2)$. These Bayes factors are shown in Fig. 4.12, calculated for a given maximum mass, i.e., the minimum speed of sound up to the corresponding M is used in the likelihood computation. There is extreme or very strong evidence against small sound speeds inside neutron stars with masses up to $M \leq 2 M_\odot$. The Bayes factors increase further with the inclusion of the black widow (BW) pulsar information. With this new data, there is strong evidence against small sound speeds, $c_{s,\min}^2 < 0.1$, inside neutron stars with masses even up to $M \leq 2.1 M_\odot$.

In Refs. [91, 133], the authors also found sound speeds larger than $c_s^2 > 0.1$ at the 95% level for neutron stars with mass $M = 2 M_\odot$. In their analyses the authors used different parametrizations and accordingly different prior distributions. With that much consistent evidence it is safe to say that a strong first-order phase transition in the core of neutron stars with mass $M \leq 2.0 M_\odot$ is fairly unlikely based on the current data. With the new information from the black widow pulsar PSR J0952-0607 the Bayes factors in Fig. 4.12 further increase such that the evidence against small sound speeds inside even heavier neutron stars with masses up to $M \leq 2.1 M_\odot$ becomes strong. The Bayes factors feature a plateau at extreme evidence for maximum masses smaller than $M \lesssim 1.9 M_\odot$, because all relevant EoSs must support these masses in order to fulfill the Shapiro and NICER constraints. Numerical values of the Bayes factors $\mathcal{B}_{c_{s,\min}^2 \leq 0.1}^{c_{s,\min}^2 > 0.1}$ for different maximum masses can be found in Table 4.4. The Bayes factors corresponding to a stronger criterion of smaller minimum speeds of sound, $\mathcal{B}_{c_{s,\min}^2 \leq 0.05}^{c_{s,\min}^2 > 0.05}$, collected in Appendix D lead to similar evidence classifications.

We find similar Bayes factors with the less general G and S versions compared in Sec. 4.1.1, see Appendix D for tabulated values. This indicates that our results are not influenced by details of the parametrization. Moreover, it turns out that just four segments are sufficient to cover the entire set of astrophysical data and theory constraints, in agreement with the findings of Refs. [91, 132]. With up to two more free segments available, the improved




FIG. 4.12. Bayes factors $\mathcal{B}_{c_{s,\min}^2 \leq 0.1}^{c_{s,\min}^2 > 0.1}$ comparing EoS samples with the following competing scenarios: a) minimum squared speed of sound (following a maximum), with $c_{s,\min}^2$ larger than 0.1, excluding a strong first-order phase transition with a Maxwell construction; versus b) EoS samples with $c_{s,\min}^2 \leq 0.1$. The Bayes factors are calculated for a given maximum neutron star mass M , i.e., the minimum speed of sound up to the corresponding mass is used. For illustration the evidence classification from Refs. [263, 264] is indicated by dashed grey lines and annotated on the right hand side. Figure taken from Ref. [97].

Segments parametrization employed in the present work has sufficient flexibility to describe additional features such as phase transitions, if these can occur within the given range of uncertainties.

In addition to setting constraints for a strong first-order phase transition in the core of neutron stars, the Bayes factors in Fig. 4.12 also limit the likelihood for the appearance of a continuous crossover with $c_s^2 \leq 0.1$. A softer crossover with $c_s^2 > 0.1$ is still possible in neutron stars with masses up to $M \lesssim 2.1 M_\odot$ and beyond. This includes EoSs featuring quark-hadron continuity [3, 235] or percolation scenarios [233], which will be discussed later in more detail. Moreover, small sound speeds $c_s^2 < 0.1$ in the cores of neutron stars with even higher masses, $M \gtrsim 2.2 M_\odot$, less constrained by the currently available astrophysical data, cannot be firmly excluded. Similarly, a phase transition with a Gibbs (rather than a Maxwell) construction [218] does not necessarily result in a drop of the sound speed to $c_s^2 \sim 0$ and can also not be ruled out.

By analyzing minima occurring at densities beyond the maximum of the speed of sound, $n(c_{s,\min}^2) > n(c_{s,\max}^2)$, we are restricting ourselves to setting constraints for strong first-order phase transitions in the deep core of neutron stars. From the behavior of the sound speed in Fig. 4.6, a minimum appearing at a density lower than that of the maximum seems to be only conceivable at small energy densities, $\varepsilon \sim 250 - 350 \text{ MeV fm}^{-3}$, i.e., at baryon densities $n \lesssim 2n_0$. With this in mind, we proceed in the next section to quantify the evidence for the possible occurrence of a phase transition strong enough to lead to a disconnected mass-radius relation.

M/M_{\odot}	$\mathcal{B}_{c_{s,\min}^2 \leq 0.1}^{c_{s,\min}^2 > 0.1}$	
	Previous	Previous + BW
1.9	201.02	500.86
2.0	46.26	229.80
2.1	4.55	15.00
2.2	1.88	3.63
2.3	1.45	2.16

TABLE 4.4. Bayes factors $\mathcal{B}_{c_{s,\min}^2 \leq 0.1}^{c_{s,\min}^2 > 0.1}$ comparing EoS samples with the following competing scenarios: a) minimum squared speed of sound (following a maximum), with $c_{s,\min}^2$ larger than 0.1, excluding a strong first-order phase transition with Maxwell construction; versus b) EoS with $c_{s,\min}^2 \leq 0.1$. The Bayes factors are calculated for a given maximum neutron star mass M , i.e., the minimum speed of sound up to the corresponding maximum mass is used in the computation.

TWIN-STAR SCENARIOS

Among the multitude of possible equations of state in the prior of the S' version, 3.5% have a disconnected mass-radius relation with more than one stable branch and hence represent a possible twin-star scenario. To quantify the evidence for such a scenario, we compute Bayes factors $\mathcal{B}_{N_{\text{branches}} > 1}^{N_{\text{branches}} = 1}$ comparing the marginalized likelihoods of EoSs with a single connected mass-radius relation to EoSs with multiple stable branches. The resulting Bayes factor of well over 900 demonstrates extreme evidence against a disconnected mass-radius relation with multiple stable branches. This value further increases with the inclusion of the new data from the black widow pulsar. The conclusion agrees well with that of Ref. [135] where the authors found only an extremely small possible parameter space for a twin-star scenario that is consistent with the low-density constraint from ChEFT and the astrophysical data. Furthermore, the authors already noted that the observation of a still more massive neutron star beyond $M \simeq 2 M_{\odot}$ would make a twin-star scenario even more unlikely. Moreover, in Ref. [105] the author also found extreme evidence that twin stars are disfavored by the current neutron star data.

If the low-density constraint involving likelihood from ChEFT is ignored, the pertinent Bayes factor decreases to $\mathcal{B}_{N_{\text{branches}} > 1}^{N_{\text{branches}} = 1} = 11.8$, providing still strong evidence against a scenario with multiple disconnected branches. In comparison, the Bayesian analyses in Refs. [74, 93], where the authors did not employ ChEFT information found only moderate evidence. This difference may be traced to the different treatment of the neutron star crust. It appears that the only possibility for a twin-star scenario, given the astrophysical database as a constraint, is through a phase transition that takes place at very low energy densities shortly above those in the neutron star crust, which was also noted in Ref. [93]. Accordingly, the mass at which the mass-radius branches become disconnected is as low as $M \sim 0.8 M_{\odot}$. (Note that in our analysis similar to Ref. [93], we do not consider disconnected branches below the assumed minimum neutron star mass, $M_{\min} = 0.5 M_{\odot}$.)

The additional inclusion of the new information from HESS J1731-347 further strengthens the evidence against a twin-star scenario, even in the absence of ChEFT constraints: the evidence now becomes very strong with a Bayes factor of $\mathcal{B}_{N_{\text{branches}} > 1}^{N_{\text{branches}} = 1} \simeq 35$. This is quite interesting as some authors considered the unusually light HESS supernova remnant as a hint in favor of a twin-star scenario [350].

MONOTONICALLY RISING SPEED OF SOUND

With priors prepared in broad generality and unrestricted initialization of the speed of sound, the previous inference results pointed out that only a weak first-order phase transition with a mixed-phase width $\Delta n/n \leq 0.2$ in a Maxwell construction can be realized inside neutron stars within the posterior credible bands of the inferred equation of state. A strong first-order phase transition is, thus, very unlikely to occur. It is worth pointing out that this conclusion is entirely based on the analysis of observational data, independent of any specific EoS model⁴. As will be discussed later, this rules out some relativistic mean-field models that suggested the occurrence of a strong (chiral) first-order phase transition already at densities below those encountered in neutron star cores. On the other hand, the empirical uncertainty bands still leave room for a continuous transition, such as a hadron-to-quark crossover [26, 232, 233, 235]. The possibility of a purely baryonic equation of state also seems feasible. One expects that such a scenario is characterized by a monotonically increasing c_s^2 as function of baryon density n , with no phase transition in the neutron star core. It is then instructive to investigate whether or not this picture is compatible with the existing empirical data.

Following Sec. 3.3.4, related insights can be gained by examining a restrictive scenario in which c_s^2 is assumed to increase monotonically up to a given transition density, n_{tr} . No phase transition or crossover occurs at baryon densities $n \leq n_{\text{tr}}$, while freedom for phase changes or the appearance of new degrees of freedom exists at densities $n > n_{\text{tr}}$, within the constraints provided by the empirical data. For the preparation of a corresponding Bayes factor analysis, we introduce a generic density, n_- , characterized by the slope $\partial c_s^2 / \partial \varepsilon < 0$ being negative at that density, i.e., the counter example to a continuously increasing sound velocity. (For n_- exceeding neutron star central densities, i.e., in the range not constrained by data, we assume that the corresponding equation of state $P(\varepsilon)$ continues rising.) Consider now the following two scenarios: hypothesis H_0 corresponds to the case $n_- > n_{\text{tr}}$, i.e., the sound velocity increases monotonically at densities up to n_{tr} . The sound speed may then change its slope and decrease at some higher density, n_- . The counter hypothesis, H_1 , assumes that this change of slope in c_s^2 occurs at a lower density instead, $n_- \leq n_{\text{tr}}$, in which case n_{tr} simply acts as a density scale for comparison with the opposite hypothesis H_0 . The Bayes factors $\mathcal{B}_{n_- > n_{\text{tr}}}^{n_- \leq n_{\text{tr}}}$ then ask for the likelihoods of the competing hypotheses and quantify the evidence of H_1 over H_0 for given values of n_{tr} . In Table 4.5 these Bayes factors are listed for different values of n_{tr} . There is strong evidence that $n_{\text{tr}} = 2n_0$ is preferred by the data. This means that an EoS with monotonically rising sound speed, $\partial c_s^2 / \partial \varepsilon > 0$ up to $n \lesssim n_{\text{tr}} = 2n_0$, is on average more likely than an EoS where the speed of sound starts to decrease in this regime, indicating that a crossover or phase transition below this transition density is unlikely. Given the previously available data there is no evidence for or against larger transition densities, which implies that a description of neutron stars in terms of nucleonic matter up to $n \sim 5n_0$, characterized by a monotonically rising speed of sound [351], cannot be excluded.

However, the inclusion of the heavy-mass measurement of the black widow pulsar PSR J0952-0607 leads to moderate evidence against a monotonically rising sound speed $\partial c_s^2 / \partial \varepsilon > 0$ up to $n \lesssim n_{\text{tr}} = 5n_0$. This means equations of state are more likely where the speed of sound drops before this transition density, hinting toward a more complex phase structure possibly with a transition to different degrees of freedom. Note that such densities

⁴ If we modify the asymptotic behavior such that the conformal limit is reached from above, there is less support for small speeds of sound, see Fig. 4.3, which makes strong first-order phase transitions inside neutron stars even more unlikely.

n_{tr}/n_0	$\mathcal{B}_{n_- > n_{\text{tr}}}^{n_- \leq n_{\text{tr}}}$	
	Previous	Previous + BW
2	0.09	0.10
3	0.40	0.59
4	1.12	1.75
5	2.11	3.16

TABLE 4.5. Bayes factors $\mathcal{B}_{n_- > n_{\text{tr}}}^{n_- \leq n_{\text{tr}}}$ comparing EoSs in which the derivative of the squared sound velocity, $\partial c_s^2/\partial\varepsilon$, turns negative at a density n_- below the transition density n_{tr} , versus EoSs with $n_- > n_{\text{tr}}$. There is strong evidence that $\partial c_s^2/\partial\varepsilon > 0$ at least up to $n_{\text{tr}} = 2n_0$. The inclusion of the heavy-mass black widow pulsar leads to moderate evidence against the speed of sound monotonically increasing up to $n_{\text{tr}} = 5n_0$.

are only reached in extremely heavy neutron stars. This evidence would further increase given a future measurement of a heavy neutron star with smaller uncertainty compared to the PSR J0952-0607 observation [96]. In this context the ongoing speculations about a very heavy neutron star based on the recent observation of gravitational wave signals from a black hole merger with a compact object of mass $2.6 M_\odot$ [352] are worth mentioning. Note, however, if the binary merger product produced in GW170817 was indeed a hyper-massive neutron star, as suggested by the measured electromagnetic counterparts [302–304], then a reduced upper limit is likely for the maximum possible mass of non-rotating neutron stars.

According to Tab. 4.5, a description of neutron star matter in terms of nucleonic degrees of freedom up to densities as high as $n_{\text{tr}} = 4n_0$ cannot be ruled out by the current astrophysical data, even with the inclusion of the heavy mass measurement of the black-widow pulsar. In Fig. 4.13 the 68% and 95% posterior credible bands of $c_s^2(\varepsilon)$ are displayed for $n_{\text{tr}} = 4n_0$. The credible bands at low energy densities $\varepsilon \lesssim 350 \text{ MeV fm}^{-3}$ differ very little from those with $n_{\text{tr}} = 0$ shown in Fig. 4.6, most likely because the data indeed prefer equations of state that rise continuously with positive curvature at least up to $\sim 2n_0$, as explained above. At higher energy densities the differences to the $n_{\text{tr}} = 0$ scenario become more apparent. Where in Fig. 4.6 the speed of sound starts forming a plateau after its rapid increase at low energy densities, for $n_{\text{tr}} = 4n_0$ the increase is less pronounced but continues up to higher energy densities, reminiscent of the posterior results for the Gaussian parametrization in Fig. 4.1. This behavior of the sound speed is also reflected in the posterior credible bands for the mass-radius relation plotted in Fig. 4.13. A restrictive prior choice of large transition densities leads to a decrease of the maximum supported mass and an increase of the central densities reached in neutron stars.

4.1.4 TRACE ANOMALY MEASURE

Based on the equation of state $P(\varepsilon)$ the trace anomaly measure Δ can be computed, given by the normalized trace of the energy momentum tensor $T^{\mu\nu}$:

$$\Delta = \frac{g_{\mu\nu} T^{\mu\nu}}{3\varepsilon} = \frac{1}{3} - \frac{P}{\varepsilon}. \quad (4.107)$$

Causality and thermodynamic stability dictate that the trace anomaly measure has to be within the range $-2/3 \leq \Delta \leq 1/3$. Moreover, $\Delta \rightarrow 0$ for conformal matter realized at high densities. The posterior credible bands for the trace anomaly measure are shown in Fig. 4.14. Starting with a value $\Delta = 1/3$ at zero density, the trace anomaly measure




FIG. 4.13. Similar to Fig. 4.6: posterior credible bands for the squared speed of sound, c_s^2 , as a function of energy density ϵ , and mass-radius relation $R(M)$ based on the previous data as well as the new information from the black-widow pulsar (BW). Here, however, for the restrictive prior case from Sec. 3.3.4 with an assumed transition density, $n_{\text{tr}} = 4n_0$, up to which the speed of sound is preconditioned to rise monotonically.

decreases with increasing energy density until $\epsilon \sim 700 \text{ MeV fm}^{-3}$, where the median of Δ turns negative. At even higher energy densities, $\epsilon \gtrsim 900 \text{ MeV fm}^{-3}$, encountered only in extremely heavy neutron stars, the 68% credible band becomes altogether negative.

In order to access the evidence for a negative trace anomaly measure we compute Bayes factors $\mathcal{B}_{\Delta \geq 0}^{\Delta < 0}$, comparing the likelihood for EoSs with negative trace anomaly, $\Delta < 0$, up to $\epsilon_{c,\text{max}}$ versus EoSs with positive Δ . Given only the previously available data with a resulting Bayes factor $\mathcal{B}_{\Delta \geq 0}^{\Delta < 0} = 6.32$, there is moderate evidence that Δ becomes negative within neutron stars. The Bayes factor further increases to $\mathcal{B}_{\Delta \geq 0}^{\Delta < 0} = 8.11$ with the inclusion of the new information from PSR J0952-0607. These results are consistent with the deduced empirical bands for Δ in Refs. [109, 206], which also start turning negative around $\epsilon \sim 700 \text{ MeV fm}^{-3}$. At the same time, the authors of Ref. [206] motivated a scenario with positive trace anomaly measure $\Delta \geq 0$, which is in light contrast to our Bayes factor analysis.




FIG. 4.14. Posterior 95% and 68% credible bands and median for the trace anomaly measure $\Delta = 1/3 - P/\epsilon$ as a function of energy density ϵ . Figure taken from Ref. [97].

Lattice QCD calculations suggest that the trace anomaly measure always stays larger than zero at finite temperatures and vanishing baryon chemical potential [209, 210]. However, in two-color QCD the trace anomaly can become negative at finite chemical potentials [203, 353]. In addition, a negative trace anomaly appears at large isospin chemical potentials

[205]. Other recent Bayesian studies have also found a negative trace anomaly measure at high baryon densities [90, 91, 94] or in extremely heavy neutron stars [134]. At much higher energy densities beyond those displayed in Fig. 4.14, the asymptotic pQCD limit does imply a switch back to positive Δ in the approach to $\Delta \rightarrow 0$.

With the density and the energy density given by

$$n = \frac{\partial P}{\partial \mu}, \quad \varepsilon = -P + \mu n, \quad (4.108)$$

the trace of the energy-momentum tensor can be expressed as

$$g_{\mu\nu}T^{\mu\nu} = \varepsilon - 3P = \mu^5 \frac{\partial}{\partial \mu} \left(\frac{P}{\mu^4} \right). \quad (4.109)$$

Thus, a negative trace anomaly measure could be a hint towards (positive) corrections to the pressure quadratic in the chemical potential, $P \propto \mu^2$. To make this more explicit, we depict the posterior credible bands for P/μ^4 as a function of μ in Fig. 4.15 and compare them to a phenomenological parametrization of the equation of state, developed originally for quark matter [354]:

$$\frac{P}{\mu^4} = -\frac{a}{\mu^4} + \frac{b}{\mu^2} + c, \quad (4.110)$$

where the coefficient a corresponds to the vacuum pressure, the quadratic coefficient b corresponds to non-perturbative power corrections to the pressure, stemming, e.g., from finite quark masses or condensates, and c corresponds to the Fermi gas contribution plus perturbative corrections, for which $P \propto \mu^4$ see Sec. 2.2.3. Fitting the phenomenological parametrization to the inferred median in Fig. 4.15 in the chemical potential range $\mu \in (950, 2000)$ MeV (at smaller chemical potentials the neutron star crust takes over) yields the following values:

$$a = 129 \text{ MeV fm}^{-3}, \quad b = 897 \text{ MeV}^2, \quad c = 1.71 \cdot 10^{-4}. \quad (4.111)$$

As displayed in Fig. 4.15, the fitted phenomenological parametrization correctly reproduces the features of the empirical pressure divided by the fourth power of the chemical potential. The value of $a^{1/4} \sim 180$ MeV is reminiscent of simplified models that describe the in-medium nucleon as a bag of quarks [355–358].

As expected from the negative trace anomaly measure, we indeed obtain a finite contribution b to the pressure quadratic in the chemical potential. Its value is much larger than the possible contributions from finite quark masses which would be proportional to $-m_f^2$. If this contribution instead results from a finite condensate, $P \propto 3/\pi^2 \mu_q^2 \Delta_{\text{gap}}^2$ [354], this would indicate a gap parameter of the order of $|\Delta_{\text{gap}}| \approx \sqrt{3\pi^2 b} \sim 160$ MeV. Such an interpretation is reminiscent of the situation at finite isospin chemical potentials, where the condensation of pions leads to a negative trace anomaly [347, 359]. At finite baryon chemical potentials the condensation of pions is unlikely [208, 360], but a recent study found that color-superconducting gaps of this magnitude are still well possible given the current neutron star data [361]. In the literature, the magnitude of the color-superconducting gap is typically given around $|\Delta_{\text{gap}}| \sim 100$ MeV [214, 215, 362]. Regarding these conclusions, it should be kept in mind that the phenomenological parametrization (4.110) was originally developed for quark matter and does not contain contributions from nucleon interactions, which will certainly be very relevant for neutron star matter.




FIG. 4.15. Posterior 95% and 68% credible bands and median for the pressure divided by the fourth power of the chemical potential, P/μ^4 , as a function of chemical potential μ . The phenomenological parametrization in Eq. (4.110) fitted to the empirical data is displayed as the dashed black line. This figure will appear in Ref. [344].

4.1.5 IMPACT OF THEORY CONSTRAINTS

IMPACT OF LOW-DENSITY ChEFT CONSTRAINT

The low-density conditioning of the EoS must incorporate the breadth of well-known empirical facts from nuclear physics. Chiral effective field theory is an established framework for this purpose. We recall that we are taking a conservative position here, terminating the applicability range of ChEFT at $n_{\text{ChEFT}} = 1.3 n_0$.

The ChEFT constraints do indeed provide an important limiting window for the evolution of the EoS to higher densities. It is thus of interest to analyze the impact of the ChEFT likelihood in the inference procedure. We do this by comparing the posterior credible bands of the speed of sound as they emerge in our approach, to the ChEFT constrained results from Refs. [60, 61] which extend up to $n_{\text{ChEFT}} = 2.0 n_0$ (see Figure 4.16). Two important findings concerning the ChEFT impact in relation to constraints from astrophysical data become apparent from this figure. First, at small energy densities the c_s^2 posterior has extra support at small sound speeds below the ChEFT constraint, because small c_s^2 are preferred by the gravitational wave event GW170817. Lowering the ChEFT constraint density to $n_{\text{ChEFT}} = 1.1 n_0$ therefore changes the posterior credible bands only marginally, see Fig. 4.17. However, at energy densities $n > n_{\text{ChEFT}} = 1.3 n_0$, the c_s^2 posterior bands increase more rapidly compared to the ChEFT constraint which remains at softer sound velocities. Accordingly, choosing $n_{\text{ChEFT}} = 2.0 n_0$ has a huge impact on the description of neutron stars: the posterior credible bands in Fig. 4.17 become more tight and the stiffening of the speed of sound is delayed to energy densities $\varepsilon \gtrsim 300 \text{ MeV fm}^{-3}$. With this change the central density of a $2.1 M_\odot$ neutron star is increased to $n_c = 3.9_{-0.8}^{+0.6} n_0$. The softening seen in the ChEFT results around $n \sim 2 n_0$ is in opposition to the apparent trend inferred from current astrophysical data. Even though there remain uncertainties about the convergence of ChEFT at higher densities, there is no indication of a steep rise in the sound speed at small densities in different ChEFT analyses [190, 363, 364]. This slight tension was already noted in Ref. [73]. There it was suggested that the range of ChEFT applicability be left as a free parameter in the range $n_{\text{ChEFT}} \sim 1.1 - 2.0 n_0$, to be sampled together with the other parameters of the EoS and then marginalized in the end [92].

It thus appears that the EoS resulting from ChEFT at densities around twice n_0 has a tendency of becoming too soft in comparison with the conditions provided by astrophysical




FIG. 4.16. Posterior credible bands for the squared speed of sound, c_s^2 , as a function of energy density ε . The inference includes previous data as well as the new information from the black-widow pulsar (BW). The low density behavior of the posterior credible bands (orange) is compared to the N^3 LO ChEFT results from Refs. [60, 61] (grey). Figure taken from Ref. [97].

data. Heavy ion collisions at intermediate energies probing the density region $n \sim 2 - 3 n_0$ may help to further clarify this situation [59, 157]. As discussed in Sec. 3.4.4, an analysis of data from the FOPI and ASY-EOS experiments led to pressure constraints similar to those derived from ChEFT [88]. On the other hand, the PREX II measurement of the ^{208}Pb neutron skin thickness suggests a stiffer EoS [82, 315]. Further insights can be expected from continuing developments in the near future.




FIG. 4.17. Posterior credible bands for the squared speed of sound, c_s^2 , as a function of energy density ε . The inference includes previous data as well as the new information from the black-widow pulsar (BW). Here, however, the ChEFT likelihood (3.89) is applied up to different densities $n_{\text{ChEFT}} = 1.1 n_0$ (left) or $2.0 n_0$ (right) instead of $n_{\text{ChEFT}} = 1.3 n_0$ suggested by Ref. [73] and used in Fig. 4.6.

IMPACT OF ASYMPTOTIC pQCD CONSTRAINT

Fig. 4.18 shows the posterior credible bands of $c_s^2(\varepsilon)$ using a different implementation of the asymptotic pQCD constraint: we shift the matching condition at which it is verified whether the asymptotic pQCD requirement can be met in a causal and thermodynamically stable fashion, from $\varepsilon_{\text{NS}} = \varepsilon_{c,\text{max}}$ to a fixed point, $n_{\text{NS}} = 10 n_0$. This is how the pQCD likelihood was implemented in Refs. [62, 86, 91]. It is obvious that this choice leads to strong, even qualitative changes in the speed-of-sound credible bands. While at energy

densities $\varepsilon \lesssim 500 \text{ MeV fm}^{-3}$, the posterior credible bands look similar to those in Fig. 4.6, at higher energy densities the speed of sound falls off and reaches significantly smaller values of c_s^2 . The Bayes factors in Table 4.5 increase, resulting in strong evidence against monotonically rising sound speeds up to densities $n_{\text{tr}} = 5 n_0$. This test case behavior agrees well with the findings in Refs. [62, 86, 91]. The differences resulting from the two pQCD implementations have already been pointed out in Refs. [93, 137].




FIG. 4.18. Similar to Fig. 4.6: posterior credible bands for the squared speed of sound, c_s^2 , as a function of energy density ε , and mass-radius relation $R(M)$ based on the previous data as well as the new information from the black-widow (BW) pulsar. Here, however, the integral pQCD likelihood is implemented, as in Refs. [62, 86, 91], at $n_{\text{NS}} = 10 n_0$ instead of $\varepsilon_{\text{NS}} = \varepsilon_{c,\text{max}}$. Figure taken from Ref. [97].

If the integral pQCD likelihood of Eq. (3.95) is imposed at densities as high as $n_{\text{NS}} = 10 n_0$, EoSs featuring large sound speeds at energy densities $\varepsilon < 1200 \text{ MeV fm}^{-3}$ are in fact not excluded, but they become much less likely. In order to fulfill the integral pQCD constraint at higher densities, EoSs with large sound speeds must decrease to smaller c_s^2 , whereas small sound speeds have more freedom to gain support in the analysis. In fact, imposing the integral pQCD likelihood at $n_{\text{NS}} = 10 n_0$ makes large speeds of sound unlikely all the way down to $\varepsilon \sim 0$. This softening does, however, depend sensitively on the specific choice of n_{NS} . For example, an alternative scenario with $n_{\text{NS}} = 8 n_0$ which is also far beyond the central densities of most neutron stars, leads to a much less pronounced softening in c_s^2 . Despite these drastic changes, in our analysis as in Refs. [86, 134], the mass-radius relation is only weakly affected. Small modifications are seen only for the most massive neutron stars, $M > 2.1 M_\odot$, which are no longer constrained by radius measurements. In fact, even the properties of a $2.3 M_\odot$ neutron star depicted in Table 4.3 change only slightly, see Table 4.6. The strong evidence against small sound speeds in the cores of neutron stars with masses $M \leq 2.0 M_\odot$ persists (see Table D.3 in Appendix D) and the maximum possible phase coexistence region within the inferred 68% posterior credible band stays below $\Delta n/n \leq 0.2$. Already in the analysis in Sec. 4.1.1 where the conformal limit is reached from above, we could conclude that the asymptotic behavior of the speed of sound has only a small influence on the properties of most neutron stars.

Because the EoSs beyond $\varepsilon_{c,\text{max}}$ are no longer constrained by astrophysical data but merely interpolated up to high densities, we believe that $\varepsilon_{\text{NS}} = \varepsilon_{c,\text{max}}$ is the better (i.e., more conservative) choice. Selecting a higher matching density may lead to an overestimation of the pQCD impact. With such a higher matching density beyond the range of control by data the impact is expected to depend sensitively on the choice of priors in the unconstrained interpolation region [93]. With the conservative choice employed in the present work, the

		Previous + BW + $n_{\text{NS}} = 10 n_0$	
		95%	68%
	n_c/n_0	$3.7^{+1.5}_{-1.1}$	$^{+0.6}_{-0.7}$
	ε_c [MeV fm $^{-3}$]	661^{+338}_{-241}	$^{+118}_{-142}$
$2.3 M_\odot$	P_c [MeV fm $^{-3}$]	223^{+190}_{-114}	$^{+61}_{-73}$
	R [km]	$12.4^{+1.1}_{-1.2}$	$^{+0.4}_{-0.7}$
	Λ	16^{+16}_{-10}	$^{+5}_{-7}$

TABLE 4.6. Same as Table 4.3: median and credible intervals for a neutron star with mass $M = 2.3 M_\odot$ are displayed given the previously available data as well as the new information from the black-widow (BW) pulsar. Here, however, the integral pQCD likelihood is implemented, as in Refs. [62, 86, 91], at $n_{\text{NS}} = 10 n_0$ instead of $\varepsilon_{\text{NS}} = \varepsilon_{c,\text{max}}$.

integral pQCD likelihood of Eq. (3.95) has a negligible influence on the sound speed and related properties of neutron stars. This corresponds well to the conclusions drawn in Ref. [137] where the authors also found only a very small impact of the pQCD integral constraint imposed at $\varepsilon_{c,\text{max}}$.

Following our results on the impact of the pQCD likelihood depending sensitively on the choice of the matching density, n_{NS} , that were published in Ref. [97], subsequent analyses found that an implementation of the pQCD integral constraint at $\varepsilon_{\text{NS}} = \varepsilon_{c,\text{max}}$ implies that the EoS has to undergo an abrupt softening after $\varepsilon_{c,\text{max}}$ to still be compatible with the pQCD results at asymptotic densities [102, 261]. Such a behavior can, of course, not be excluded a priori. This observation led to the development of a new pQCD likelihood, that marginalizes over a Gaussian process additionally conditioned on the sound speed computed in pQCD [261]. In this way, the likelihood becomes independent of the specific choice of n_{NS} at the cost of an increased model dependence on the hyper-parameters of the Gaussian process. Regarding these questions, we want to reiterate the small impact seen of the more restrictive pQCD likelihood at $n_{\text{NS}} = 10 n_0$ even on heavy neutron stars with a mass as high as $2.3 M_\odot$ as noted in Table 4.6. Thus, in all cases the pQCD likelihood only affects the properties of extremely heavy neutron stars not constrained by any current astrophysical data.

Note that perturbative QCD calculations generally assume a trivial unpaired ground state. However, as explained in Sec. 2.2.4, due to the attractive gluonic force between quarks a color superconductor with quark Cooper pairs is expected to be the true ground state at asymptotically large chemical potentials. Recent analyses [215, 365] found that the resulting finite gap in the quark excitation spectrum could increase the sound velocity and pressure at large densities beyond the values for unpaired quark matter, reminiscent of the impact of pion condensation at large isospin chemical potentials [204, 205, 347].

Some authors recently claimed evidence for a possible phase transition to a new state of matter near the maximum neutron star mass, M_{max} [91, 92], based on inferred values for the sound speed [95, 134], the polytrope index [130] or the behavior of the trace anomaly measure at densities corresponding to M_{max} . In these analyses the authors use the approach with fixed matching density to implement the asymptotic pQCD constraint, which, as we discussed, is strongly prior dependent. In addition, the central densities reached inside the most massive neutron stars with masses close to M_{max} are, similar to our analysis, much higher than the central density of a $2.1 M_\odot$ neutron star. This high density regime is,

however, only loosely constrained by the current astrophysical data. Therefore, analyses of properties of the most massive stars should require a particularly detailed assessment of the prior dependence induced by the interpolation to high densities. In contrast, in our analysis we claim evidence only for neutron stars with masses $\lesssim 2.1 M_\odot$, which are still in the density regime that is well-constrained by the current astrophysical data.

A further note concerns selected recent analyses which saw a pronounced softening of the sound velocity and equation of state at high energy densities [90–92, 95, 132]. Apart from a different implementation of the pQCD likelihood, as discussed above, this may be due to differences in presentation. As pointed out, in computing credibility bands we follow Refs. [67, 70] and employ each equation of state only up to their respective maximum central energy density $\varepsilon_{c,\max}$ corresponding to the respective maximum mass. Similarly, in the computation of the credible bands for the radius as a function of mass we consider each EoS only up to their respective maximum mass M_{\max} , as higher masses do not correspond to stable neutron stars. In contrast, Refs. [90–92, 95, 132] among others use each EoS up to arbitrarily high energy densities. We decided against this approach because it generates an uncontrolled mix of information from EoSs constrained by astrophysical data and EoSs beyond empirical limits. If we also use each equation of state up to arbitrary ε , we do indeed observe a slight softening in Fig. 4.6. However, this behavior is not based on the empirical data but merely on the interpolation extending up to the pQCD constraint at asymptotic densities.

4.1.6 FURTHER MEASUREMENTS

POSSIBLE IMPACT OF HESS J1731-347

Next, we analyze the impact of including the mass-radius estimate for the very light central compact object HESS J1731-347 reported in Ref. [45]⁵. The updated posterior credible bands including this additional information are collected in Fig. 4.19. To reach small radii, the speed of sound has to increase more rapidly at densities above the ChEFT constraint at $n_{\text{ChEFT}} = 1.3 n_0$. With the new information, the credible bands are much more tightly constrained compared to Fig. 4.6. The inclusion of the supernova remnant in the Bayesian analysis shifts the radii at all masses to lower values. The radius of a $1.4 M_\odot$ neutron star reduces to $R = 11.8^{+0.5}_{-0.4}$ km at the 68% level, similar to the value $R = 11.7 \pm 0.5$ km reported in Ref. [45] at the 90% level. In addition to HESS J1731-347 and the previously available data listed in Table 3.3, the latter estimate includes additional information from the X-ray burster 4U 1702-429 and from the rotation limit for the radio pulsar PSR J1748-2446ad. There is visibly some tension between the radius estimate at lower masses based on the current data, most importantly from PSR J0030+0451, and HESS J1731-347, as already noted in Ref. [95, 103]. As a consequence of this tension the posterior credible band for $R(M)$ agrees with the credible intervals of the supernova remnant only marginally at the 95% level. However, as also noted in Sec. 3.4.2, the analysis of HESS J1731-347 involves more systematic uncertainties compared to other data including NICER.

In Table 4.7 we show inferred properties of neutron stars with mass $M = 0.77 M_\odot$. The central density of a $0.77 M_\odot$ neutron star is low, only $n_c = 2.2 \pm 0.3 n_0$. As an additional output, the Bayes factor $\mathcal{B}_{\Delta \geq 0}^{\Delta < 0}$ with the inclusion of the HESS supernova remnant increases from about 8 to $\mathcal{B}_{\Delta \geq 0}^{\Delta < 0} \simeq 11$, so that the evidence for a negative trace anomaly measure inside neutron stars turns from moderate to strong.

⁵ Note that due to the inclusion of a neutron star crust in our EoS parametrizations, we cannot analyze a hypothetical strange star scenario in our analysis [366].




FIG. 4.19. Similar to Fig. 4.6: posterior credible bands for the squared speed of sound, c_s^2 , as a function of energy density ε , and mass-radius relation $R(M)$, but now including the information from the supernova remnant HESS J1731-34 in addition to the black-widow (BW) pulsar PSR J0952-0607. The resulting mass-radius relation is also compared to the marginalized intervals at the 68% level from the analysis of HESS J1731-34 [45] (green). Figure taken from Ref. [97].

	Previous + BW + HESS	
	95%	68%
n_c/n_0	$2.2^{+0.6}_{-0.5}$	± 0.3
ε_c [MeV fm $^{-3}$]	352^{+81}_{-87}	$^{+38}_{-52}$
$0.77 M_\odot$ P_c [MeV fm $^{-3}$]	25^{+9}_{-7}	± 4
R [km]	$11.8^{+0.9}_{-1.0}$	$^{+0.5}_{-0.4}$
Λ	7911^{+3979}_{-3352}	$^{+1799}_{-1827}$

TABLE 4.7. Same as Table 4.2, but median and credible intervals for a neutron star with mass $M = 0.77 M_\odot$ are displayed, given the previously available data, the mass measurement of the black widow pulsar (BW) and the mass-radius data of the supernova remnant HESS J1731-347.

IMPACT OF PSR J0437-4715

Finally, we examine the implications of incorporating the preliminary mass-radius estimate for PSR J0437-4715 obtained from the newest NICER (NN) measurement. Fig. 4.20 displays the updated posterior credible bands for the sound velocity and the mass-radius relation. With the integration of the new mass-radius data, which has the lowest uncertainty of all NICER data to date, the posterior bands are more tightly constrained. Similar to the impact of the supernova remnant HESS J1731-347 depicted in Fig. 4.19, the sound speed must increase more rapidly to accommodate for the pulsar’s small radius. Consequently, this shifts the mass-radius relation towards lower radii, leading to median values around ~ 12.0 km, similar to the inference results seen in Refs. [106, 107].

Table 4.8 collects properties of neutron stars with characteristic masses $M = 1.4 M_\odot$ and $2.1 M_\odot$. The radius and tidal deformability of a $1.4 M_\odot$ neutron star reported in Ref. [107] align closely with our results. When compared to the values listed in Table 4.2, it becomes apparent that the latest NICER measurement not only reduces the estimated radii of neutron stars, but also leads to higher internal central densities. Nevertheless, even at a mass of $2.3 M_\odot$ the central densities remain below five times nuclear saturation density




FIG. 4.20. Similar to Fig. 4.6: posterior credible bands for the squared speed of sound, c_s^2 , as a function of energy density ε , and mass-radius relation $R(M)$, but now including the information from the newest NICER (NN) measurement of PSR J0437-4715 in addition to the black-widow (BW) pulsar PSR J0952-0607. The resulting mass-radius relation is also compared to the marginalized intervals at the 68% level from the preliminary analysis of PSR J0437-4715 [46] (red).

at the 68% level, $n_c(2.3, M_\odot) = 4.1^{+0.7}_{-0.8} n_0$. Similar to the low-radius constraint from HESS J1731-347, including the information from the newest NICER measurement leads to strong evidence for a negative trace anomaly measure inside neutron stars, possibly caused by a finite condensate as discussed in Sec. 4.1.4. Future data from the NICER telescope will help to refine the posterior credible bands even further.

		Previous + BW + NN	
		95%	68%
	n_c/n_0	2.8 ± 0.7	$^{+0.3}_{-0.4}$
	ε_c [MeV fm $^{-3}$]	454^{+117}_{-119}	$^{+58}_{-66}$
$1.4 M_\odot$	P_c [MeV fm $^{-3}$]	68^{+28}_{-23}	$^{+11}_{-16}$
	R [km]	12.0 ± 0.9	$^{+0.4}_{-0.5}$
	Λ	356^{+223}_{-150}	$^{+80}_{-114}$
	n_c/n_0	$3.9^{+1.6}_{-1.4}$	$^{+0.6}_{-0.8}$
	ε_c [MeV fm $^{-3}$]	677^{+364}_{-271}	$^{+121}_{-174}$
$2.1 M_\odot$	P_c [MeV fm $^{-3}$]	215^{+195}_{-122}	$^{+57}_{-84}$
	R [km]	$11.8^{+1.2}_{-1.1}$	± 0.6
	Λ	21^{+24}_{-16}	$^{+8}_{-11}$

TABLE 4.8. Similar to Table 4.2: median and credible intervals for neutron stars with masses $M = 1.4 M_\odot$ and 2.1 are displayed, but now including the information from the newest NICER (NN) measurement of PSR J0437-4715 in addition to the previously available data and the black widow (BW) pulsar PSR J0952-0607.

4.2 PHENOMENOLOGY AND MODELS

The previous sections have shown how present and forthcoming neutron star data (masses, radii, and tidal deformabilities) provide constraints for the equation of state and for possible phase transitions in cold, dense baryonic matter. The enhanced stiffness of the EoS at high baryon densities is a necessary feature; sufficiently high pressures are required to support neutron stars with masses around and beyond $2 M_\odot$ at relatively large radii ($R \simeq 11 - 13$ km). The central core densities are not as extreme as was previously imagined. Even in the core of a $2.3 M_\odot$ neutron star, the baryon density at its center stays below five times the equilibrium density of nuclear matter (at the 68% credible level; see Table 4.3). The upper border of the 95% credible interval of the inferred central density $n_c(2.3 M_\odot) = 3.8_{-1.3}^{+1.6} n_0$ lies slightly above $5 n_0$, but with a very low probability. Very similar credible intervals are found for the G and S versions compared in Sec. 4.1.1. This observation has consequences for the interpretation of the possible structure and composition of neutron star matter, which we discuss in the following⁶.

4.2.1 LOW-ENERGY NUCLEON STRUCTURE

Spontaneously broken chiral symmetry, as the long-wavelength manifestation of QCD, governs the low-energy structure and dynamics of hadrons, including nucleons and pions. As chiral Nambu–Goldstone bosons, pions play a distinguished role in this context. Models based on chiral symmetry often view the nucleon as a complex system of two scales [367]: a compact hard core that hosts the three valence quarks and, thus, encloses the baryon number and a surrounding quark-antiquark cloud in which pions figure prominently as the ‘soft’ degrees of freedom.

Such a two-scale scenario is manifested in empirical form factors and sizes of nucleons [163, 368]. Consider, for example, the proton and neutron electromagnetic form factors and their slopes at zero momentum transfer, which determine the corresponding mean square radii. The empirical r.m.s. proton charge radius, $\langle r_p^2 \rangle^{1/2} = 0.840 \pm 0.003 \pm 0.002$ fm [369], combined with the neutron charge radius, $\langle r_n^2 \rangle = -0.105 \pm 0.006$ fm² [370], gives the isoscalar and isovector mean square radii of the nucleon, $\langle r_{S,V}^2 \rangle = \langle r_p^2 \rangle \pm \langle r_n^2 \rangle$, with the following resulting values:

$$\sqrt{\langle r_S^2 \rangle} \simeq 0.78 \text{ fm} , \quad \sqrt{\langle r_V^2 \rangle} \simeq 0.90 \text{ fm} . \quad (4.112)$$

Each of the nucleon form factors $G_i(q^2)$ related to a current with index i has a representation in terms of a once-subtracted dispersion relation,

$$G_i(q^2) = G_i(0) + \frac{q^2}{\pi} \int_{t_0}^{\infty} dt \frac{\text{Im} G_i(t)}{t(t - q^2 - i\epsilon)} , \quad (4.113)$$

with the squared four-momentum transfer $q^2 = q_0^2 - \vec{q}^2$. The mean square radii are given as

$$\langle r_i^2 \rangle = \frac{6}{G_i(0)} \left. \frac{dG_i(q^2)}{dq^2} \right|_{q^2=0} = \frac{6}{\pi} \int_{t_0}^{\infty} \frac{dt}{t^2} S_i(t) , \quad (4.114)$$

where the distribution $S_i(t) = \text{Im}G_i(t)/G_i(0)$ represents the spectrum of intermediate hadronic states through which the external probing field couples to the respective nucleon current. For example, the isovector charge radius reflects the interacting two-pion cloud

⁶ This section closely follows Ref. [163].

of the nucleon governed by the ρ meson and a low-mass tail extending down to the $\pi\pi$ threshold, $t_0 = 4m_\pi^2$. The isoscalar charge radius is related to the three-pion spectrum, which is strongly dominated by the narrow ω meson [371] and starts at $t_0 = 9m_\pi^2$. The isoscalar charge form factor of the nucleon, $G_S^E(q^2)$ (with $G_S^E(0) = 1$), is particularly suitable for discussing a delineation between the ‘core’ and ‘cloud’ parts of the nucleon [372, 373]. The vector meson dominance principle implies, in its simplest version, a representation of the form

$$G_S^E(q^2) = \frac{F_B(q^2)}{1 + |q^2|/m_\omega^2} . \quad (4.115)$$

The form factor $F_B(q^2)$ of the baryon number distribution in the nucleon core acts as a source for the ω field that propagates with its mass m_ω . Introducing the mean square radius of the baryon core, $\langle r_B^2 \rangle = 6 \frac{dF_B(q^2)}{dq^2} |_{q^2=0}$, the mean square isoscalar charge radius becomes

$$\langle r_S^2 \rangle = \langle r_B^2 \rangle + \frac{6}{m_\omega^2} . \quad (4.116)$$

Using $m_\omega = 783$ MeV, the estimated baryonic core radius is

$$\sqrt{\langle r_B^2 \rangle} \simeq 0.47 \text{ fm} . \quad (4.117)$$

A nucleon core size of about 1/2 fm is characteristic of chiral ‘core + cloud’ models. It also holds up in more detailed treatments of the spectral distributions governing the nucleon form factors [369]. The inclusion of additional ϕ meson and $\pi\rho$ continuum contributions in the spectral function of $G_S^E(q^2)$ moves the core radius to just slightly larger values [368].

Consider as another example the form factor associated with the axial vector current of the nucleon. The corresponding mean-square axial radius deduced from neutrino-deuteron scattering data is reported as [374]

$$\langle r_A^2 \rangle = (0.46 \pm 0.22) \text{ fm}^2 . \quad (4.118)$$

A schematic axial vector dominance picture would assign a dominant part of the ‘cloud’ contribution to this form factor through the spectrum of the $a_1(1260)$ meson with its large width. If an approximate scale of this ‘cloud’ part is identified with $\delta\langle r_A^2 \rangle \sim 6/m_a^2$ using the physical a_1 mass, $m_a \simeq 1.23$ GeV, one finds

$$\sqrt{\langle r_A^2 \rangle}_{\text{core}} = [\langle r_A^2 \rangle - \delta\langle r_A^2 \rangle]^{1/2} \simeq 0.55 \text{ fm} , \quad (4.119)$$

with an estimated uncertainty of about 25%.

Yet another interesting piece of information is the mass radius of the proton deduced from J/Ψ photoproduction data [375]. It involves the form factor of the trace, T_μ^μ , of the nucleon’s energy-momentum tensor and is supposed to be dominated by gluon dynamics at the center of the nucleon. Low-mass (e.g., two-pion and -kaon) components are suppressed [368]. The result quoted in [375],

$$\langle r^2 \rangle_{\text{mass}}^{1/2} = (0.55 \pm 0.03) \text{ fm} , \quad (4.120)$$

is, once again, remarkably close to an assumed ‘core’ size scale of $\sim 1/2$ fm.

These empirical considerations motivate a picture of the nucleon as containing a compact ‘hard’ core in which the valence quarks and their baryon number are confined, and a ‘soft’ surface of quark-antiquark pairs forming a mesonic cloud. This structure has implications

for the behavior of nucleons in dense baryonic matter. With a core size of $R_{\text{core}} \sim 1/2$ fm and a cloud range given, e.g., by the proton charge radius, $R_{\text{cloud}} \sim 0.84$ fm, there is a significant separation of volume scales in vacuum: $(R_{\text{cloud}}/R_{\text{core}})^3 \sim 5$.

This scale separation is expected to increase further in dense baryonic matter for the following reasons. The properties of the soft multi-pion cloud are closely tied to spontaneously broken chiral symmetry and the approximate Nambu–Goldstone boson nature of the pion. The size of this cloud is expected to increase with density, along with the decreasing in-medium pion decay constant, $f_{\pi}^*(n)$, which acts as a chiral order parameter. The baryonic core, on the other hand, is governed by gluon dynamics, without a leading connection to chiral symmetry in QCD. This core is, therefore, expected to be quite stable against changes from increasing density up until the compact hard cores begin to touch and overlap. What arises in this way is a two-scale scenario for dense baryonic




FIG. 4.21. Sketch of low- and high-density baryonic matter. Baryons (e.g., nucleons) are viewed as valence quark cores surrounded by clouds of quark-antiquark pairs (e.g., chiral meson clouds). At densities of $n \gtrsim 2 - 3 n_0$, the percolation of quark-antiquark pairs over larger distances starts, as indicated. Valence quark cores begin to touch and overlap at baryon densities $n \gtrsim 5 n_0$. Figure taken from Ref. [163].

matter, as described in [233] and sketched in Figure 4.21. At $n \simeq n_0$, the tails of the meson clouds of nucleon pairs overlap, resulting in two-body exchange forces. As the average distance between nucleons decreases with increasing density, around $n \gtrsim 2 - 3 n_0$, the soft clouds of $q\bar{q}$ pairs start to be delocalized. Their mobility expands over larger distances in a percolation process involving larger numbers of nucleons. In the terms of conventional nuclear physics, this corresponds to the emergence of many-body forces, the strength of which grows with increasing density. At that stage, the baryonic cores are still separated but subject to increasingly repulsive Pauli principle effects. The cores begin to touch and overlap at average nucleon-nucleon distances of $d \lesssim 1$ fm, corresponding to densities $n \gtrsim 5 n_0$. Further compression of baryonic matter would still have to overcome the strongly repulsive hard core in the nucleon-nucleon interaction. Recalling the inferred credible intervals of the central densities in heavy neutron stars in Table 4.3, one concludes that a phase transition to valence quark deconfinement does not seem likely in a two-scale picture and under the conditions in the cores of neutron stars unless they are extremely heavy.

An interesting and closely related result emerges from a detailed analysis of y -scaling in electron-nucleus scattering at large momentum transfers ($|\vec{q}| \gtrsim 1$ GeV) and low energy transfers [376]. The persistently observed y -scaling under these kinematical conditions implies that the electrons still scatter from strongly correlated pairs of nucleons, rather than quarks, at short distances. The conclusion drawn in [376] is that at local densities as large as

five times n_0 , nuclear matter still appears to behave as a collection of nucleons. As discussed in Sec. 4.1.3, such a picture is still well compatible with the data set consisting of Shapiro time delay, NICER, ChEFT, pQCD and gravitational wave data. On the other hand, the heavy-mass measurement of the black-widow pulsar hints towards a non-monotonic behavior of the speed of sound at densities $n > 4 n_0$ (see Tab. 4.5).

4.2.2 QUARK-HADRON CONTINUITY

While a strong first-order phase transition in neutron star matter seems unlikely based on the current empirical observations, a continuous crossover from hadrons to quarks is still possible within the present data-driven constraints. Such a scenario is realized, for example, in the QHC21 equation of state [235]. It features a smooth interpolation between low and high densities from nuclear to quark matter regimes. The quark matter phase is described by a three-flavor Nambu–Jona-Lasinio (NJL) model that includes pairing (diquark) degrees of freedom and a strongly repulsive vector coupling between quarks. In order to build up the necessary pressure at high densities, this vector coupling must be comparable to or larger than the strength of the standard scalar-pseudoscalar interaction in the NJL model. The density at which the interpolated turnover to quark matter takes place in the QHC21 EoS is chosen as $n \gtrsim 3.5 n_0$, within the range of central densities that can be reached in $M \gtrsim 2 M_\odot$ neutron stars. A crossover from hadrons to quarks could also involve an intermediate phase of quarkyonic matter as explained in Sec. 2.2.4.

4.2.3 CHIRAL SYMMETRY RESTORATION

The quest for chiral symmetry restoration at high baryon densities – a transition from the spontaneously broken Nambu–Goldstone realization to the unbroken Wigner–Weyl mode – has been a persistent issue for a long time. A possible first-order chiral phase transition and the existence of a corresponding critical endpoint in the QCD phase diagram have always been topics of prime interest [377–379]. As discussed in Sec. 2.2.4, early hypotheses concerning the occurrence of a first-order phase transition were frequently based on Nambu–Jona-Lasinio (NJL)-type models in mean-field approximation [380–382], which were later extended by incorporating confinement aspects through added Polyakov-loop degrees of freedom [383–385].

The empirical constraints on strong first-order phase transitions in dense neutron star matter, as described in Sec. 4.1.3, are strikingly at variance with previous mean-field (e.g., NJL model) predictions. These suggested that a first-order chiral phase transition should already appear at relatively low baryon densities around $n \sim 2 - 3 n_0$. A possible explanation for this discrepancy can be found in [220, 221], where a chiral nucleon-meson (ChNM) field theory was used to explore the effects of fluctuations beyond mean-field (MF) approximation. The starting point was a relativistic chiral Lagrangian, $\mathcal{L}(\Psi; \pi, \sigma; v_\mu)$, shaped around a linear sigma model with nucleons (Ψ), pions (π), and a scalar (σ) field. Short-range dynamics were parametrized in terms of heavy isoscalar and isovector vector fields (v_μ). The expectation value $\langle \sigma \rangle$ of the scalar field acted as a chiral order parameter normalized in the vacuum to the pion decay constant, $f_\pi \simeq 92$ MeV. Two classes of fluctuations beyond MF were then systematically studied: first, vacuum fluctuations that introduced an additional term proportional to $\sigma^4 \ln(\sigma/f_\pi)$ in the MF partition function; secondly, fluctuations involving pion loops and nucleon particle-hole excitations. The vacuum fluctuations can be included in an extended mean-field (EMF) approximation [386]. Fluctuations involving pion and nucleon loops are computed using non-perturbative func-

tional renormalization group (FRG) methods. The parameters of the ChNM model – in particular, those related to short-distance dynamics – are fixed to reproduce empirical nuclear physics data [220, 221].

Fig. 4.22 demonstrates the important role of fluctuations beyond the mean-field approximation for the chiral order parameter $\langle\sigma\rangle$. In symmetric nuclear matter, the mean-field approximation of the ChNM model correctly reproduces the first-order liquid-gas phase transition at low density. However, the chiral order parameter also displays a strong first-order chiral phase transition with a Maxwell-constructed phase coexistence region starting already below $2n_0$. For neutron matter, which has no liquid-gas phase transition, the MF approximation nevertheless predicts a first-order chiral phase transition at densities around $n \sim 3n_0$, which is well within the range of densities realized in neutron stars. However, in both nuclear and neutron matter, the inclusion of fermionic vacuum fluctuations (i.e., the effect of the ground state zero-point energy) in the extended mean-field (EMF) approximation converts the first-order chiral phase transition into a smooth crossover and shifts it to densities of $n > 5n_0$. This effect is further enhanced by the additional fluctuations included in the full FRG calculation, as demonstrated in Fig. 4.22. As a result, the restoration of chiral symmetry is relegated to very high baryon densities beyond the inferred core densities in even the heaviest neutron stars (see Table 4.3). A comparable impact of fluctuations on the phase structure is seen in alternative chiral models [387, 388].

Another approach based on chiral symmetry is the parity-doublet model. In this model, the active coupled baryonic degrees of freedom are the nucleon with spin-parity $1/2^+$ and its chiral partner with spin-parity $1/2^-$, where the latter is identified with the $N^*(1535)$ resonance. Spontaneous chiral symmetry breaking in vacuum manifests itself in the mass splitting of these two states, while the in-medium restoration pattern of this symmetry is signaled by the N and N^* masses becoming degenerate. A recent detailed analysis [389] of the chiral order parameter in this model using extended mean-field approximation found a chiral phase transition in nuclear matter, but at extremely high densities ($n > 10n_0$) that are far beyond the density scales reached in neutron stars.

4.2.4 FERMI LIQUID PICTURE

According to Sec. 4.1.3, the constraints on the equation of state inferred from the current empirical data still permit an interpretation of neutron star core matter in terms of baryonic degrees of freedom, such as a system dominated by neutrons [268] with small fractions of protons and perhaps hyperons [238, 240, 390]. The inferred baryon chemical potential $\mu(n) = \partial\varepsilon(n)/\partial n$ shown in Fig. 4.9 does not distinguish between different species of baryons. Its behavior nonetheless displays increasingly strong correlations at high densities. It is instructive to analyze the gross properties of this state of matter using the Landau theory of relativistic Fermi liquids [391], which describes a relativistic many-body system of strongly interacting fermions at low temperatures. In this theory, the ground state is represented by a filled Fermi sea of quasiparticles. These quasiparticles can be understood as free fermions dressed by the strong interactions with the surrounding medium. Quasiparticles above the Fermi surface and quasiholes below it correspond to low-lying excitations [392]. Here, we perform a schematic study assuming ‘neutron-like’ quasiparticles [351] with a baryon number of $B = 1$ and a density-dependent mass of $m(n)$ while ignoring other small admixtures in the composition of the dense medium. The quasiparticles are characterized by their (relativistic) Landau effective mass m_L^* at a Fermi momentum of $p_F = (3\pi^2 n)^{1/3}$,

$$m_L^*(n) = \sqrt{p_F^2 + m^2(n)}, \quad (4.121)$$




FIG. 4.22. Chiral order parameters in symmetric nuclear matter and neutron matter at a temperature of $T = 0$ as a function of baryon density in units of nuclear ground state equilibrium density, $n_0 = 0.16 \text{ fm}^{-3}$. Dotted lines: liquid-gas phase transition (L-G) in symmetric nuclear matter. Dashed lines: first-order chiral phase transitions emerging from the mean-field (MF) approximation of a relativistic chiral nucleon-meson (ChNM) field-theoretical model. Solid lines show the results of extended mean-field (EMF) calculations (with the inclusion of fermionic vacuum fluctuations) and functional renormalization group (FRG) calculations based on the same ChNM model. Figures taken from Ref. [163], where they were adapted from Refs. [220, 221].

together with an effective potential, $U(n)$, so that the baryon chemical potential can be written as [351]

$$\mu(n) = m_L^*(n) + U(n) . \quad (4.122)$$

The median of the $\mu(n)$ posterior credible bands in Fig. 4.9 is now taken as a guiding starting point to extract the baryonic quasiparticle properties. With an educated ansatz for $m(n)$, the density dependence of the potential $U(n)$ can then be deduced and further discussed. One possible choice is to take $m(n)$ (with $m(0) = 939 \text{ MeV}$) from the non-perturbative FRG calculation employing the chiral nucleon-meson field-theoretical model [220] discussed in Sec. 4.2.1. The resulting Landau effective mass $m_L^*(n)$ is shown in Fig. 4.23 together with the potential $U(n) = \mu(n) - m_L^*(n)$. The 95% credible band of $\mu(n)$ in Fig. 4.9 leads to an uncertainty of about 15% for U at high densities. It is instructive to fit the resulting quasiparticle potential by a series in powers of baryon density for $n \lesssim 5 n_0$ (with $n_0 = 0.16 \text{ fm}^{-3}$, as before):

$$U(n) = \sum_n u_n \left(\frac{n}{n_0} \right)^n . \quad (4.123)$$

The coefficients fitted to the median of $\mu(n)$ are

$$u_1 = 90.9 \text{ MeV}, \quad u_2 = 15.3 \text{ MeV}, \quad u_3 = 3.2 \text{ MeV}, \quad u_4 = -0.4 \text{ MeV} . \quad (4.124)$$

This pattern reflects a hierarchy of many-body correlations, recalling that the term linear in density represents two-body interactions, the term of order n^2 corresponds to short-range

three-body forces, and so forth. The role of the repulsive N -body terms with $N > 2$ is quite significant; at $n \simeq 4n_0$, corresponding to an average distance of about 1 fm between the baryonic quasiparticles, these terms contribute as much as the two-body forces to the potential U and generate the strong pressure to support heavy neutron stars. Of course, these statements rely on the ansatz for the density-dependent mass $m(n)$, which is guided by the FRG calculations of pure neutron matter. A small fraction of protons in beta-equilibrated matter will not substantially change this picture. However, neutron star core compositions that qualitatively deviate from this simplified picture may lead to different conclusions.

Finally, consider the dimensionless Fermi liquid parameters, F_0 and F_1 , of the spin-averaged quasiparticle interaction. In terms of the quasiparticle mass and potential, they are given as [351]

$$F_0(n) = \frac{p_F}{\pi^2} \left[m(n) \frac{\partial m}{\partial n} + m_L^*(n) \frac{\partial U}{\partial n} \right], \quad F_1(n) = -\frac{3U(n)}{\mu(n)}. \quad (4.125)$$

Further useful relations are $1 + F_0 = N(0)(\partial\mu/\partial n)$ with the density of quasiparticle states at the Fermi surface, $N(0) = m_L^* p_F / \pi^2$, and $1 + F_1/3 = 1 - U/\mu = m_L^*/\mu$.




FIG. 4.23. The Landau effective mass $m_L^*(n) = \sqrt{p_F^2 + m^2(n)}$ and potential $U(n)$ of quasiparticles representing the median of the posterior distribution of the baryon chemical potential $\mu(n)$ (see Fig. 4.9). Figure taken from Ref. [163].

The result in Fig. 4.24 shows a strongly increasing F_0 at high densities. This reflects once again the growing importance of many-body correlations as matter becomes more and more compact. Such repulsive correlations are responsible for the increase in the sound velocity beyond its canonical conformal limit, $c_s^2 > 1/3$, as seen in Fig. 4.6. The Fermi liquid parameter F_1 is smaller in magnitude and has a negative slope, indicating the decreasing effective mass at higher densities. The Landau parameter F_0 displays the typical behavior of a strongly correlated Fermi system. However, it is interesting to observe that, in comparison with the leading Landau parameters in liquid ${}^3\text{He}$, the correlations in neutron star core matter are not extraordinarily strong. Values of $F_0 \approx 9.3$ were reported for ${}^3\text{He}$ at zero pressure, and $F_0 \approx 68.2$ at a pressure of 27 bar [393]. Accordingly, the Fermi liquid parameters for matter in the density range that is reached inside even the heaviest neutron stars, $n \lesssim 5n_0$, are significantly smaller. From this perspective, the dense baryonic medium encountered in the center of neutron stars is perhaps not as extreme as often imagined.




FIG. 4.24. The Landau Fermi liquid parameters derived from quasiparticle properties that are based on the median of the data-inferred baryon chemical potential μ (see Fig. 4.9). Figure taken from Ref. [163].

4.3 INFERENCE DIRECTLY FROM TELESCOPE SPECTRA

In the previous sections, we presented and discussed posterior results obtained using the traditional two-step inference approach, i.e, based on a set of exterior neutron star properties that were inferred from observed detector data in a previous step. With the simulation-based inference method introduced in Sec. 3.5, inference of the full posterior distribution of neutron star EoS and nuisance parameters directly from telescope spectra of quiescent low-mass binaries is now feasible⁷. Following Refs. [127, 128], 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 in Sec. 3.5.1. 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 neural likelihood estimation 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.

4.3.1 EXAMPLE POSTERIOR DISTRIBUTION

An example of the posterior distributions marginalized to one and two dimensions is shown in the corner plot in Fig. 4.25. 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 $\log(T_{\text{eff}}^{(1)})$. While the marginalized posteriors of the nuisance parameters for the other nine stars are also available, they are not shown here. In all three nuisance parameter scenarios from Table 3.5, the EoS parameters are strongly correlated, similar to the log-likelihoods computed in Ref. [128]. 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 specified in Sec. 3.3.1.

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

⁷ This section closely follows Ref. [243].

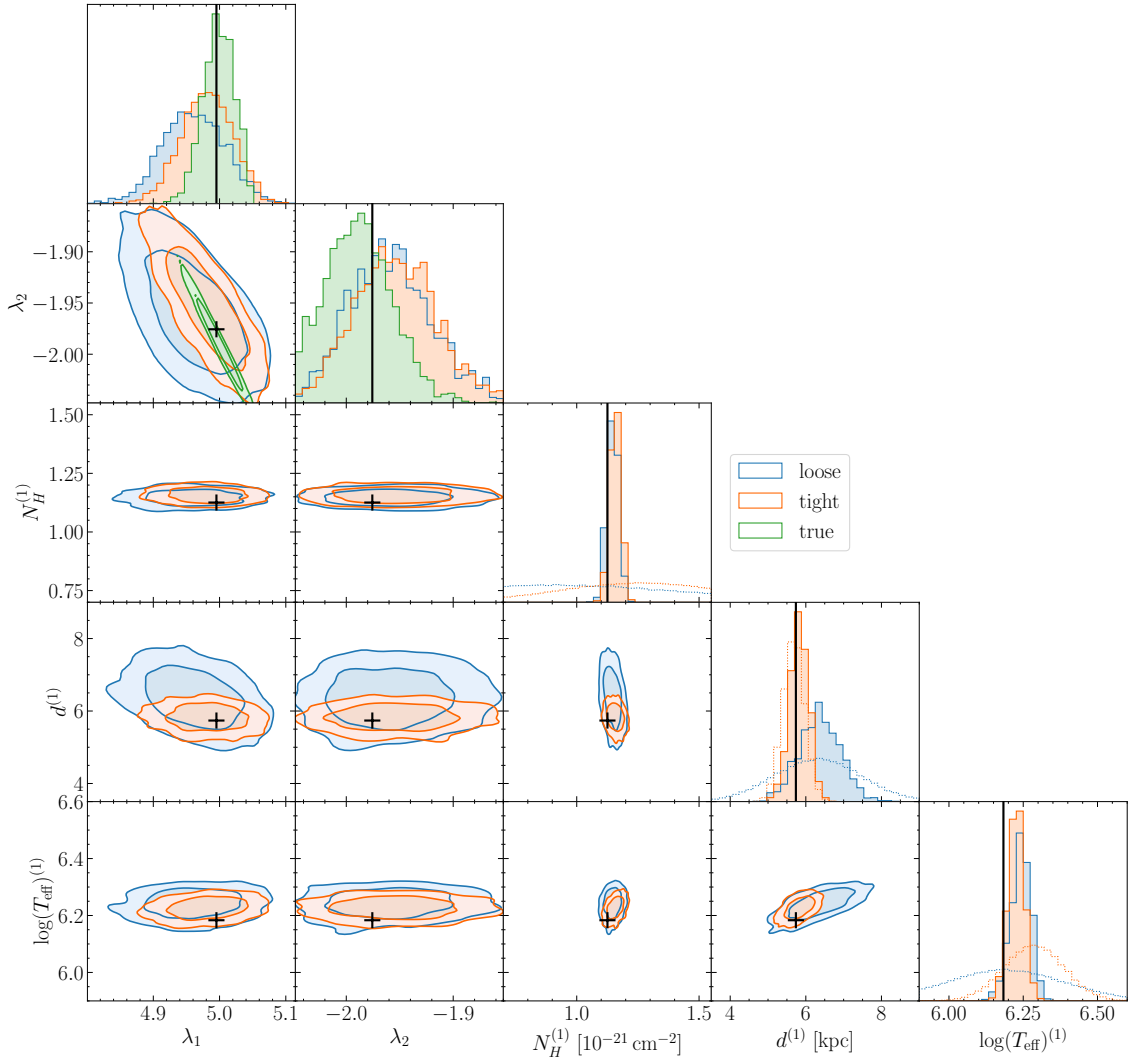


FIG. 4.25. Corner plot depicting the posterior distribution of the parameters λ_1 and λ_2 of one example EoS as well as the first 3 (of 30) nuisance parameters $N_H^{(1)}$, $d^{(1)}$ and $\log(T_{\text{eff}})^{(1)}$. The posterior is computed based on the simulated spectra of 10 stars with the nuisance parameters known exactly in the true scenario (green), and known with the uncertainties in Table 3.5 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. Figure taken from Ref. [243].

distributions for the EoS parameters. This is further pronounced for the loose case, where less prior information on the nuisance parameters is available. Fig. 4.25 illustrates that the hydrogen column N_H as well as the logarithm of the effective surface temperature $\log(T_{\text{eff}})$ can be significantly constrained from the spectrum data compared to their prior ranges. In the tight scenario, the marginal posterior for the distance d is almost indistinguishable from the prior, indicating that the telescope spectra do not contribute any more information for this parameter over the tight priors. However, in the loose case, the marginal posterior distribution of d becomes tighter than the loose prior, which implies that we can indeed extract information about the distance of a neutron star from its X-ray spectrum⁸.

⁸ The same holds true for the nuisance parameters of the other nine neutron stars not shown here.

We can transform the posterior distribution for the EoS parameters λ_1 and λ_2 into 95% posterior credible bands for the pressure as a function of energy density as depicted in Fig. 4.26. As before, the constraints are much tighter in the true scenario and become increasingly broader in the tight and loose cases. By solving the TOV equations, we can translate the EoS into mass-radius constraints. The 95% posterior credible bands for the radius as a function of mass are depicted in Fig. 4.26. Focusing only on the tight case in Fig. 4.27, 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.




FIG. 4.26. Posterior 95% credible bands for the pressure as a function of energy density and the radius as a function of mass for the three (true, tight, loose) scenarios. Similar to Fig. 4.25, 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 taken from Ref. [243].




FIG. 4.27. Similar to Fig. 4.26, 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. Figure taken from Ref. [243].

4.3.2 INCREASING THE NUMBER OF OBSERVATIONS

With an anticipated surge in the number of available neutron star observations in the upcoming years (see Sec. 3.4), the inference method must be able to scale to a large set of data. In our novel approach, normalizing flows approximate the likelihood $p(s|\theta, \nu)$ per observed star. These likelihoods are then combined to obtain the total likelihood for a set of neutron stars, see Eq. (3.102). Consequently, the method does not need to assume a fixed number of observed neutron stars, nor a particular ordering of the stars, and works out-of-the-box for a growing set of observed neutron stars without the need for retraining any networks. To demonstrate this, we present the marginal posterior distributions for one example EoS in Fig. 4.28, for 5, 10, and 20 observed neutron stars, and with loose uncertainties on the nuisance parameters.




FIG. 4.28. 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 3.5 of the *loose* scenario. The ground-truth parameter values are depicted as black crosses/lines. Figure taken from Ref. [243].

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 HMC for posterior sampling, the computation time is predominantly consumed by the evaluation of the likelihood and its gradient. While the availability of more observations increases the per-iteration computational time in sampling the posterior, it also speeds up the convergence of the algorithm. Furthermore, HMC algorithms can easily be scaled to thousands of dimensions, hence we do not anticipate the dimensionality to be a limiting factor in the scaling of our approach.

4.3.3 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 machine learning (ML) approaches that infer the neutron star EoS directly from telescope spectra described in Sec. 3.5.2, we use the same accuracy measure as Refs. [127, 128]. For each EoS in the test set, we simulate 10 spectra with random nuisance parameters. Based on the spectra and the prior nuisance parameter information, we then sample the posterior using the methodology outlined in Sec. 3.5.3. From the marginal posterior distributions, similar to the example in Fig. 4.25, we determine the maximum-a-posteriori estimates (MAP)⁹ for the two EoS parameters ($\lambda_{1,\text{pred}}, \lambda_{2,\text{pred}}$) and compare them to the ground-truth values ($\lambda_{1,\text{truth}}, \lambda_{2,\text{truth}}$). The distributions of the differences between the marginal MAP estimates and the ground-truth values, ($\lambda_{1,\text{pred}} - \lambda_{1,\text{truth}}, \lambda_{2,\text{pred}} - \lambda_{2,\text{truth}}$), on the test data set are depicted in Fig. 4.29. 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 around zero, indicating that there is no systematic bias in the posterior prediction. Compared to the previous ML analyses in Refs. [127, 128], the distributions of λ_1 are much narrower in the tight and loose scenarios (see for example Fig. 8 in [128]).




FIG. 4.29. 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 3.5. Figure taken from Ref. [243].

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 Fig. 4.29. We can combine both standard deviations into

$$\sigma_{\text{tot}} = \sqrt{\sigma_{\lambda_1}^2 + \sigma_{\lambda_2}^2}. \quad (4.126)$$

The resulting values are listed in Table 4.9 and illustrated in Fig. 4.30. 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 3.5. In the true case, the NLE approach performs better than ML-Likelihood_{EoS} from Ref. [128] and NN(Spectra) from Ref. [127] (see Sec. 3.5.2 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

⁹ Note that this is the maximum-a-posteriori estimate of the marginal and not of the full posterior.

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 4.9 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 equation of state 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.

$p(\nu)$	Method	$\lambda_{1,\text{pred}} - \lambda_{1,\text{truth}}$		$\lambda_{2,\text{pred}} - \lambda_{2,\text{truth}}$		Combined
		μ	σ	μ	σ	σ_{tot}
true	ML-Likelihood _{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 _{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 _{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 4.9. 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 Fig. 4.29, i.e., of the differences between the predicted MAP and ground-truth values. Both standard deviations are combined to σ_{tot} according to Eq. (4.126). 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. [127], as well as an approach using an approximate likelihood that incorporates two neural networks, ML-Likelihood_{EoS}, from [128]. 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 3.5.

With our neural likelihood estimation approach, we can now, for the first time, infer the posterior for nuisance parameters. The means and standard deviations of the differences between the marginal MAP estimates and the ground-truth values on the test data are listed in Table 4.10. While we obtain excellent constraints on the hydrogen fraction N_H and the effective surface temperature $\log(T_{\text{eff}})$, our ability to constrain the distance d from the spectra is limited. When going from the tight to the loose scenario, the inference on the distance estimates suffers the most, indicating that these constraints are primarily driven by the prior in the tight case, as seen in Fig. 4.25. 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.

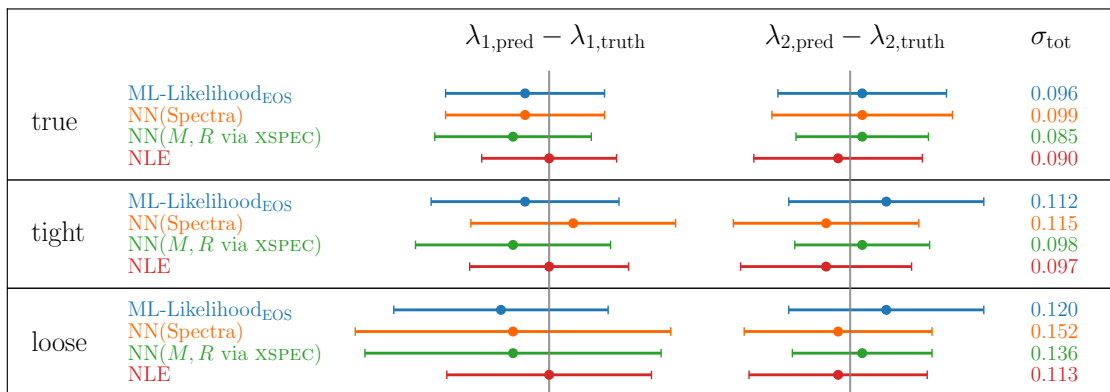


FIG. 4.30. Illustrated mean and standard deviation of the difference between the predicted maximum-a-posteriori values to the ground-truth values for the three different scenarios from Table 4.9. Figure taken from Ref. [243].

Method	$p(\nu)$	$N_{H,\text{pred}} - N_{H,\text{truth}}$		$d_{\text{pred}} - d_{\text{truth}}$		$\log(T_{\text{eff}})_{\text{pred}} - \log(T_{\text{eff}})_{\text{truth}}$	
		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 4.10. Similar to Table 4.9, but with mean and standard deviation of the predicted maximum-a-posteriori estimates versus the ground-truth values for the three nuisance parameters.

4.3.4 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 [287] 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 Sec. 4.3.2. In addition, extending to additional stars is inexpensive relative to other methods, which require integrating over estimated mass-radius posteriors to construct likelihoods [97, 103], such as with kernel density estimation techniques (see, e.g., Eq. (3.71)).
- (iii) Learning the likelihood instead of the posterior allows combination with likelihoods from other data [394], e.g., the neutron star data introduced in Sec. 3.4. 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.
- (v) These techniques can be extended to NICER or gravitational wave data [395]¹⁰. Many more gravitational wave events from binary neutron star mergers will be detected in future runs of LIGO, Virgo, and KAGRA. In addition, the gravitational wave database will increase dramatically with the next generation of gravitational wave detectors as explained in Sec. 3.4. For such a large number of expected observations, it is not numerically feasible to compute the likelihood by integrating KDEs as in Eqs. (3.71) and (3.87) [120, 394, 396]. Neural simulation-based inference techniques promise to provide a cost-efficient alternative [328, 397].
- (vi) This approach can help guide decisions about future observations. Sec. 4.3.2 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.

Before applying neural likelihood estimation to real observed telescope spectra 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 $p(s|\theta, \nu, M)$ rather than $p(s|\theta, \nu)$, 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., [41, 45, 141, 398, 399]. Preliminary results show, however, that the mass distributions can become multimodal, as also seen in [141], 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 parametrizations of the EoS, such as the Segments parametrization introduced in Sec. 3.1.2, which can describe more complicated phase structures possibly realized inside neutron stars like crossovers of first-order phase transitions. Our novel simulation-based inference approach is conducive to the use of a more complex EoS parametrization with more parameters.

¹⁰For the analysis of NICER and gravitational waves, the detector data is much larger compared to the qLMXBs analyzed here. In that case, it might become necessary to use an additional embedding net to learn summary statistics [338].

SUMMARY AND CONCLUSIONS

In this thesis, we have used Bayesian inference methods to translate empirical multimessenger data of neutron stars into constraints on the neutron star matter equation of state (EoS). Here, the likelihoods for the astrophysical detector data are analytically unavailable. Therefore, we first followed previous work to infer the EoS from a set of masses, radii and tidal deformabilities previously extracted from observed detector data.

A primary aim is to tighten the conclusions about possible phase transitions in the cores of neutron stars. A key quantity to address these issues is the squared speed of sound, $c_s^2 = \partial P / \partial \varepsilon$. We have modeled c_s^2 using two generic parametrizations, a skewed Gaussian in combination with a logistic function, and a more general form based on piecewise linear segments. We used a multimessenger data set consisting of Shapiro time delay observations of selected pulsars, NICER X-ray measurements and gravitational wave signals from binary neutron star mergers. The asymptotic behavior of the squared sound speed is matched to perturbative QCD calculations. At low baryon densities, nuclear physics constraints from chiral effective field theory, treated as a likelihood rather than a prior, are incorporated up to a conservative window of baryon densities of $n \leq 1.3 n_0$, with $n_0 = 0.16 \text{ fm}^{-3}$ the equilibrium density of normal nuclear matter.

Good agreement is found between the output posteriors for both parametrizations, as well as for versions with a changed asymptotic behavior and an increased number of free segments. Nevertheless, the Segments parametrization leads to slightly larger posterior credible bands due to its less restrictive functional form. The overall conclusion is that the Bayesian inference approach generates results that are stable with respect to variations in the functional form of the prior if the initial parametrization is chosen sufficiently general. In the energy density range $\varepsilon \sim 0.3 - 0.5 \text{ GeV fm}^{-3}$ the squared speed of sound rises rapidly beyond $c_s^2 = 1/3$ and develops a plateau at larger energy densities. Moderate tension exists between ChEFT extrapolations of c_s^2 up to $n \simeq 2 n_0$ and the trend towards a stiffer EoS implied by the astrophysical data. Matching the neutron star EoS to asymptotic pQCD requires an extrapolation from densities reached in neutron star cores, n_{NS} , to the extreme densities at which perturbative QCD methods can be applied. If instead a matching density is chosen at a value far beyond the density range controlled by empirical observations, the impact of the pQCD constraint on the EoS for neutron stars depends sensitively on this prior choice. In all cases, we find only very little impact of the asymptotic pQCD constraint on the properties of neutron stars with masses $M \lesssim 2.3 M_\odot$. Furthermore, a Bayes factor analysis suggests moderate evidence for a negative trace anomaly measure, $\Delta = 1/3 - P/\varepsilon$, in heavy neutron stars ($M \gtrsim 2 M_\odot$). The evidence increases to strong after including preliminary results from the newest NICER measurement. This could be a hint for a possible finite condensate at high baryon densities, reminiscent of the situation at finite isospin densities.

The incorporation of the heaviest neutron star observed so far, the black widow pulsar PSR J0952-0607, with rotation corrections properly applied, results in a further stiffening of the EoS. This leads to the remarkable feature that the median of the mass-radius relation is at an almost constant radius ($R \simeq 12.3$ km) for all masses above $M \geq 0.7 M_\odot$. As a consequence of the stiffer EoS, the central core densities of neutron stars are reduced to $n_c < 5 n_0$ even for masses as high as $M \simeq 2.3 M_\odot$ (at 68% credibility). In a baryon-dominated system, such densities correspond to average distances between any two baryons of 1 fm or larger. Chiral model descriptions of the nucleon suggest a scale separation between a compact valence quark core and a mesonic cloud, so the nucleonic cores with radii of about 1/2 fm would begin to touch and overlap only in the deep interior of extremely heavy neutron stars. The appearance of a strong first-order phase transition becomes unlikely under such conditions. This is consistent with the empirical results: within the inferred 68% posterior credible bands, possible phase coexistence regions, i.e., domains of constant pressure in a Maxwell construction, are restricted to a maximal width of $\Delta n/n \leq (\Delta n/n)_{\max} \simeq 0.2$. Moreover, the inclusion of the heavy black-widow pulsar increases the evidence against small sound speeds inside neutron star cores. A corresponding Bayes factor investigation demonstrates strong evidence against minimum squared sound speeds smaller than $c_{s,\min}^2 \leq 0.1$, indicative of a possible strong first-order phase transition, in neutron stars with masses up to $M \leq 2.1 M_\odot$. On the other hand, a continuous hadron-to-quark crossover scenario – or alternatively, baryonic matter as a strongly correlated relativistic Fermi liquid – remains possible within the inferred data-driven constraints.

Finally, we have developed a novel approach based on recently established simulation-based inference methods that can infer the EoS parameters directly from astrophysical detector data in a single step. In neural likelihood estimation (NLE), normalizing flows are trained on samples of simulated detector observations to learn the analytically unavailable likelihoods of the neutron star data. This allows the usage of Hamiltonian Monte Carlo (HMC) methods, which use derivatives to sample more efficiently from the resulting posterior distribution. We have inferred the parameters of a spectral EoS parametrization based on simulated thermal X-ray spectra of quiescent low-mass binaries. The telescope spectra of these emissions depend on the mass and radius of the star and are also affected by stellar nuisance parameters, which are not always well-constrained. Using NLE in conjunction with HMC methods, we can, for the first time, infer the full posterior distribution of both the equation of state and nuisance parameters directly from telescope observations. In realistic scenarios for nuisance parameter priors, this method outperforms all state-of-the-art methods in terms of prediction accuracy. Our approach improves the interpretability of the results by providing access to the full posterior distribution, and naturally scales to the growing number of neutron star observations expected from the next generation of X-ray and gravitational wave detectors. As the number of available neutron star data increases, conventional Bayesian methods with their expensive integrations become impractical, whereas neural likelihood estimation techniques could provide a computationally efficient alternative.

With an increase of the quantity and quality of observational data, coupled with progress in more efficient inference methods, we are looking forward to more rigorous constraints on the EoS and the phase structure of dense matter in the future.



EFFECTIVE ENTHALPY

Instead of numerically solving the Tolman–Oppenheimer–Volkoff equations (2.8) and (2.9) for $P(r)$ and $m(r)$, following Ref. [166] we can use a different approach based on the effective enthalpy defined as:

$$h(P) = \int_0^P \frac{dP'}{\varepsilon(P') + P'} . \quad (\text{A.127})$$

In neutron stars $h(r)$ is monotonically decreasing from the center until it vanishes at the star's surface. For a given equation of state, $P(\varepsilon)$, pressure and energy density can be determined in terms of the effective enthalpy, i.e., $P(h)$ and $\varepsilon(h)$, using the above expression. With the derivative of h given by

$$\frac{\partial h}{\partial r} = \frac{\partial h}{\partial P} \frac{\partial P}{\partial r} = - \frac{m + 4\pi r^3 P(h)}{r(r - 2m)} , \quad (\text{A.128})$$

we can rewrite the TOV equations in terms of the effective enthalpy:

$$\frac{\partial m}{\partial h} = - \frac{4\pi \varepsilon(h) r^3 (r - 2m)}{m + 4\pi r^3 P(h)} , \quad (\text{A.129})$$

$$\frac{\partial r}{\partial h} = - \frac{r(r - 2m)}{m + 4\pi r^3 P(h)} . \quad (\text{A.130})$$

This set of differential equations can be numerically integrated for a given central enthalpy $h(r = 0) = h_c$. The total mass and radius of the star can then be determined from $M = m(h = 0)$ and $R = r(h = 0)$. Note that in the star's center where r and m equal zero, the above differential equations become singular. Therefore, we have to determine the boundary conditions for a given central enthalpy from a series expansion for $r(h)$ and $m(h)$ near $h = h_c$ [400]

$$r(h) = r_1 (h_c - h)^{1/2} + r_3 (h_c - h)^{3/2} + \mathcal{O}(h_c - h)^{5/2} , \quad (\text{A.131})$$

$$m(h) = m_3 (h_c - h)^{3/2} + m_5 (h_c - h)^{5/2} + \mathcal{O}(h_c - h)^{7/2} , \quad (\text{A.132})$$

with the following coefficients:

$$\begin{aligned} r_1 &= \left[\frac{3}{2\pi(\varepsilon_c + 3P_c)} \right]^{1/2} , & r_3 &= - \frac{r_1}{4(\varepsilon_c + 3P_c)} \left[\varepsilon_c - 3P_c - \frac{3}{5} \left(\frac{\partial \varepsilon}{\partial h} \right)_c \right] , \\ m_3 &= \frac{4\pi}{3} \varepsilon_c r_1^3 , & m_5 &= 4\pi r_1^3 \left[\frac{r_3 \varepsilon_c}{r_1} - \frac{1}{5} \left(\frac{\partial \varepsilon}{\partial h} \right)_c \right] , \end{aligned} \quad (\text{A.133})$$

and the central energy density and pressure given by $\varepsilon_c = \varepsilon(h_c)$ and $P_c = P(h_c)$. Solving Eqs. (A.129) and (A.130) with the above boundary conditions for different central enthalpies h_c yields the mass-radius relation (M, R) for a given neutron star EoS. Similarly, the differential equations for the tidal deformability in Sec. 2.1.3 can be reexpressed in terms of the effective enthalpy [400].

B

ROTATION CORRECTION

In Sec. 2.1.2, we briefly reviewed the derivation the TOV equations from the theory of general relativity for spherically symmetric, static sources. More complex differential equations can be derived for uniformly rotating neutron stars [401]. They are, however, too numerically expensive to be employed in Bayesian inference analyses which require mass-radius sequences for a large amount of equations of state. It is therefore beneficial to instead use universal relations that relate properties of static stars to their rotating counterparts. Here, we summarize the empirical formulas derived in Ref. [272] to determine the mass M_{rot} and equatorial radius R_{rot} of a star uniformly rotating with an angular velocity $\Omega = 2\pi\nu$ from the mass M and radius R of a non-rotating star.

In their work, the authors first derived an empirical formula to determine the angular velocity of a star spinning at the Kepler limit, i.e., the maximum spin-rate for a neutron star before mass-shedding at the equator occurs (see Eq. (2.2)), based on the dimensionless compactness of the non-rotating star, $C = M/R \text{ km}/M_{\odot}$:

$$\Omega_K = \Omega_* \sum_{i=0}^4 a_i C^i, \quad (\text{B.134})$$

with the normalization constant $\Omega_* = \sqrt{G_N M/R^3}$. The authors found that the mass and radius including rotation effects can be determined from the dimensionless compactness C and the angular velocity normalized to the Kepler frequency $\Omega_n = \Omega/\Omega_K$:

$$\frac{M_{\text{rot}}}{M} = 1 + \left(e^{A_m \Omega_n^2} - 1 \right) \left(\sum_{i=0}^4 a_{m,i} C^i \right), \quad (\text{B.135})$$

$$\frac{R_{\text{rot}}}{R} = 1 + \left(e^{A_r \Omega_n^2} - 1 + B_r \left[\log \left(1 - \left(\frac{\Omega_n}{1.1} \right)^4 \right) \right]^2 \right) \left(1 + \sum_{i=1}^5 a_{r,i} C^i \right). \quad (\text{B.136})$$

These empirical formulas were fitted to a set of equations of state sampled from a piecewise polytrope and a skewed Gaussian parametrization. The resulting values for the parameters $(A_m, A_r, B_r, a_i, a_{m,i}, a_{r,i})$ are listed in Table B.1. For stars rotating slower than the Kepler limit the authors found deviations of less than 5% for the mass and radius of a rotating star determined using the above empirical formulas. This is much smaller than the uncertainty of the mass measurement of PSR J0952-0607, $M = 2.35 \pm 0.17 M_{\odot}$, to which the universal formulas are applied.

	a_0	a_1	a_2	a_3	a_4	a_5	A	B
Ω_K	0.552	3.304	-35.211	180.61	-326.48			
R_{rot}/R		-15.496	442.60	-4945.62	23458.06	-40544.25	0.203	0.1611
M_{rot}/M	-0.0160	3.123	-20.721	41.202	-6.464		1.127	

TABLE B.1. Empirical coefficients for the formulas (B.134 - B.136) derived in Ref. [272] to compute the Kepler frequency, Ω_K , as well as the mass, M_{rot} , and radius, R_{rot} , of rotating neutron stars.

EOS TABLES

For practical purposes and applications, the median values of the squared sound velocity as a function of energy density, as shown in Fig. 4.4, are listed in Table C.1, using the ‘Previous’ data set consisting of the Shapiro time delay, NICER, ChEFT, pQCD and gravitational wave data¹. Based on these values the pressure P is computed using Eq. (2.11), as well as the baryon density n and the chemical potential μ with Eq. (4.106). The asymmetry of the posterior distribution causes small insignificant deviations between the pressure computed from the integral of the sound velocity and the median of P in Fig. 4.4. Similarly, the median values of $c_s^2(\varepsilon)$ including the new information from the black-widow pulsar depicted in Fig. 4.6 are listed in Table C.2. A comparison between both tables is instructive as it underlines the significant stiffening of the EoS that emerges when the new PSR J0952-0607 data are incorporated.

ε [GeV fm ⁻³]	Previous			
	c_s^2	P [MeV fm ⁻³]	n/n_0	μ [GeV]
0.1	0.02	0.8	0.66	0.96
0.2	0.07	5.5	1.31	0.98
0.3	0.18	17.4	1.93	1.03
0.4	0.36	43.8	2.51	1.11
0.5	0.50	87.7	3.05	1.20
0.6	0.57	142.0	3.55	1.31
0.7	0.60	200.7	4.01	1.40
0.8	0.60	260.3	4.45	1.49
0.9	0.59	319.5	4.86	1.57
1.0	0.58	377.7	5.24	1.64
1.1	0.58	435.5	5.61	1.71
1.2	0.59	493.8	5.97	1.77

TABLE C.1. Tabulated values of the median for the squared sound velocity, c_s^2 , as a function of energy density ε as shown in Fig. 4.4, i.e., inferred from the Shapiro time delay, NICER, ChEFT, pQCD and gravitational wave data listed as ‘Previous’ in Table 3.3. Based on these values the pressure is computed as well as the baryon density n (in units of the nuclear saturation density n_0) and the baryon chemical potential μ .

¹ Parts of the text in this appendix have been adapted from Ref. [97].

ε [GeV fm ⁻³]	Previous + BW			
	c_s^2	P [MeV fm ⁻³]	n/n_0	μ [GeV]
0.1	0.02	0.8	0.66	0.96
0.2	0.07	5.4	1.31	0.98
0.3	0.20	17.4	1.93	1.03
0.4	0.44	49.0	2.50	1.12
0.5	0.59	101.3	3.03	1.24
0.6	0.64	163.0	3.52	1.35
0.7	0.64	226.8	3.97	1.46
0.8	0.62	289.8	4.39	1.55
0.9	0.60	350.9	4.78	1.64
1.0	0.59	410.6	5.15	1.71
1.1	0.60	470.4	5.50	1.78
1.2	0.61	531.1	5.84	1.85

TABLE C.2. Similar to Table C.1: tabulated values of the median for $c_s^2(\varepsilon)$, but now including the new information from PSR J0952-0607 in addition to the previously available data as depicted in Fig. 4.6. Based on the median we again compute the pressure P , density n and chemical potential μ .

D

BAYES FACTORS

A key result of the present work is the systematics of the Bayes factor $\mathcal{B}_{c_{s,\min}^2 \leq 0.1}^{c_{s,\min}^2 > 0.1}$, quantifying the evidence against a dropping of the squared sound speed to values below $c_{s,\min}^2 \leq 0.1$, as a function of the maximum mass in neutron stars displayed in Fig. 4.12 and Table 4.4¹. For a further documentation of these results tables of numerical values for additional scenarios are useful to underscore the increase of this evidence with the inclusion of the new information from the heavy mass measurement of PSR J0952-0607.

First, in Table D.1 we display the Bayes factors for the Gaussian and ‘basic’ version of the Segments parametrization with fixed asymptotic behavior. For both parametrizations we find even larger Bayes factors compared to the S’ version discussed in the main text. Nevertheless, the Bayes factors result in a similar evidence classification, i.e., there is strong evidence that $c_{s,\min}^2$ does not become smaller than 0.1 in neutron stars with mass $M \leq 2.0 M_\odot$. With the new empirical information from the black widow pulsar the Bayes factors further increase and there is strong evidence against small sound speeds $c_{s,\min}^2 < 0.1$ inside neutron stars even up to masses $M \leq 2.1 M_\odot$ for both parametrizations. This illustrates that the constraints on strong first-order phase transitions are driven by the data and do not depend on the chosen functional form.

M/M_\odot	$\mathcal{B}_{c_{s,\min}^2 \leq 0.1}^{c_{s,\min}^2 > 0.1}$			
	Previous		Previous + BW	
	Gaussian	Segments	Gaussian	Segments
1.9	1.03×10^4	2.05×10^5	4.13×10^4	6.71×10^5
2.0	74.06	74.26	388.78	317.08
2.1	6.12	4.67	16.32	18.65
2.2	3.08	2.32	4.95	5.96
2.3	2.70	1.67	3.42	2.51

TABLE D.1. Bayes factors $\mathcal{B}_{c_{s,\min}^2 \leq 0.1}^{c_{s,\min}^2 > 0.1}$ comparing EoS samples with the following competing scenarios: a) minimum squared speed of sound (following a maximum), with $c_{s,\min}^2$ larger than 0.1, excluding a first-order phase transition; versus b) EoS with $c_{s,\min}^2 \leq 0.1$. The Bayes factors are calculated for a given maximum neutron star mass M , i.e., the minimum speed of sound up to the corresponding mass is used.

¹ Parts of the text in this appendix have been adapted from Ref. [97].

In Table D.2 we display Bayes factors corresponding to a stronger criterion of smaller minimum speeds of sound, $c_{s,\min}^2 \leq 0.05$, for the improved Segments parametrization. Although there are some minor modifications compared to the Bayes factors that consider minimum sound speeds smaller or equal to $c_{s,\min}^2 \leq 0.1$ listed in Table 4.4, both cases yield the same evidence classification against small sound velocities in the cores of neutron stars.

M/M_\odot	$\mathcal{B}_{c_{s,\min}^2 \leq 0.05}^{c_{s,\min}^2 > 0.05}$	
	Previous	Previous + BW
1.9	259.81	416.14
2.0	48.86	220.60
2.1	4.50	14.47
2.2	2.13	4.42
2.3	1.52	2.20

TABLE D.2. Similar to Table 4.4: Bayes factors $\mathcal{B}_{c_{s,\min}^2 \leq 0.05}^{c_{s,\min}^2 > 0.05}$, quantifying the evidence against a dropping of the squared sound speed to values below $c_{s,\min}^2 \leq 0.05$ in neutron stars, as a function of the maximum mass.

Finally, in Table D.3 we display the Bayes factors $\mathcal{B}_{c_{s,\min}^2 \leq 0.1}^{c_{s,\min}^2 > 0.1}$ using a different implementation of the asymptotic pQCD constraint, i.e., shifting the matching condition from $\varepsilon_{\text{NS}} = \varepsilon_{c,\max}$ to $n_{\text{NS}} = 10 n_0$ as in Refs. [62, 86, 91]. With this changed asymptotic behavior the Bayes factors decrease, but there is still strong evidence against small minimum sound speeds in the cores of neutron stars with masses up to $\leq 2.0 M_\odot$, and the Bayes factors increase again significantly with the inclusion of the heavy-mass measurement. Despite the softening seen in the speed of sound at high energy densities in Fig. 4.18, there is no significant preference for small $c_{s,\min}^2$ even in neutron stars as heavy as $2.3 M_\odot$.

M/M_\odot	$\mathcal{B}_{c_{s,\min}^2 \leq 0.1}^{c_{s,\min}^2 > 0.1}$	
	Previous + $n_{\text{NS}} = 10 n_0$	Previous + BW + $n_{\text{NS}} = 10 n_0$
1.9	101.01	211.62
2.0	23.45	93.49
2.1	2.16	5.82
2.2	0.84	1.32
2.3	0.64	0.76

TABLE D.3. Similar to Table 4.4: Bayes factors $\mathcal{B}_{c_{s,\min}^2 \leq 0.1}^{c_{s,\min}^2 > 0.1}$, quantifying the evidence against a dropping of the squared sound speed to values below $c_{s,\min}^2 \leq 0.1$ in neutron stars, as a function of the maximum mass. Here, however, the integral pQCD likelihood is implemented, as in Refs. [62, 86, 91], at $n_{\text{NS}} = 10 n_0$ instead of $\varepsilon_{\text{NS}} = \varepsilon_{c,\max}$.

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 unnormalized densities¹.

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., $y \sim \mathcal{N}(x, \sigma)$. The proposal is then accepted with some probability $\alpha(x, y)$, in which case the next state is y , otherwise it remains x . We need to identify this acceptance probability to ensure that asymptotically this Markov chain generates samples from the target distribution, $x, y \sim \pi$. This is equivalent to ensuring that the target distribution π is invariant under this Markov chain, for example by maintaining detailed balance conditions when accepting the proposed stage:

$$\pi(x)t(x, y)\alpha(x, y) = \pi(y)t(y, x)\alpha(y, x) . \quad (\text{E.137})$$

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

If $t(x, y)$ is positive everywhere, the above condition is satisfied by the standard Metropolis–Hastings acceptance formula for $x, y \in \mathcal{R}^n$,

$$\alpha(x, y) = \min\left(\frac{\pi(y)t(y, x)}{\pi(x)t(x, y)}, 1\right), \quad (\text{E.138})$$

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. (E.137), but with lower acceptance rates, meaning they lead to an undesirable higher asymptotic variance for estimated expectations.

¹ This appendix closely follows Ref. [243].

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 [147, 148]. 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 $H : \mathcal{R}^{2n} \rightarrow \mathcal{R}$ gives the total energy for the state $x := (q, p)$,

$$H(x) = H(q, p) = U(q) + \frac{1}{2}p^T M^{-1}p . \quad (\text{E.139})$$

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

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

$$\tilde{\pi}(x) := Z^{-1} e^{-H(x)} = Z^{-1} e^{-U(q)} e^{-\frac{1}{2}p^T M^{-1}p} = Z^{-1} \pi(q) e^{-\frac{1}{2}p^T M^{-1}p} , \quad (\text{E.140})$$

where $Z = \int_{\mathcal{R}^{2n}} dx e^{-H(x)} = (2\pi)^{n/2} \sqrt{\det M}$ is a normalizing constant. Since H is the sum of potential and kinetic terms, in the Gibbs density q and p are independent, with the q -marginal of $\tilde{\pi}(x)$ being the target density $\pi(q)$. Thus, extracting the first n coordinates of samples $x^{(i)}$ from $\tilde{\pi}$, one obtains samples from 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. (E.139) over a certain time T , starting from the initial point $x = (q^{(i)}, p^{(i)})$ and is followed by negation of the final momentum². This map defines the transition kernel t . This proposal is then accepted with the probability $\alpha(x, y)$.

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

$$\begin{aligned} \bar{p} &\leftarrow p - \frac{\epsilon}{2} \nabla U(q) , \\ q' &\leftarrow q + \epsilon M^{-1} \bar{p} , \\ p' &\leftarrow \bar{p} - \frac{\epsilon}{2} \nabla U(q') . \end{aligned} \quad (\text{E.141})$$

² 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.

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 $P(q, p) := (q, -p)$, $F = (L_\epsilon)^n P$ (recalling that operators act left to right) is volume-preserving because each of the three substeps is a shear transformation³, and F is an involution⁴ because $L_\epsilon(q', -p') = (q, -p)$, which can be verified by reversing the order of the substeps, so $L_\epsilon P L_\epsilon = P$ and $((L_\epsilon)^n P)^2 = I$. In this case, the transition described in Step 2 preserves detailed balance when

$$\alpha(x, y) = \min \left(\frac{\tilde{\pi}(y)}{\tilde{\pi}(x)}, 1 \right), \quad (\text{E.142})$$

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.

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 [402]. We use an importance sampling step to determine initial values for each chain [326].

³ 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.

⁴ Involution means $F^{-1} = F$, i.e., it is *time reversible*

BIBLIOGRAPHY

- [1] L. McLerran, *Acta Phys. Pol. B* **51**, 1067 (2020).
- [2] F. Özel and P. Freire, *Ann. Rev. Astron. Astr.* **54**, 401–440 (2016).
- [3] G. Baym, T. Hatsuda, T. Kojo, P. D. Powell, Y. Song, and T. Takatsuka, *Rep. Prog. Phys.* **81**, 056902 (2018).
- [4] I. Vidaña, *Eur. Phys. J. Plus* **133**, 445 (2018).
- [5] J. Nättilä and J. J. E. Kajava, in *Handbook of X-ray and Gamma-ray Astrophysics* (Springer Nature Singapore, Nov. 2022), pp. 1–53.
- [6] N. Yunes, M. C. Miller, and K. Yagi, *Nature Rev. Phys.* **4**, 237–246 (2022).
- [7] L. Landau, *Physikalische Zeitschrift der Sowjetunion* **1**, 285–288 (1932).
- [8] W. Baade and F. Zwicky, *Phys. Rev.* **46**, 76–77 (1934).
- [9] R. C. Tolman, *Phys. Rev.* **55**, 364–373 (1939).
- [10] J. R. Oppenheimer and G. M. Volkoff, *Phys. Rev.* **55**, 374–381 (1939).
- [11] A. Hewish, S. J. Bell, J. D. H. Pilington, P. F. Scott, and R. A. Collins, *Nature* **217**, 709–713 (1968).
- [12] D. Blaschke and N. Chamel, in *The Physics and Astrophysics of Neutron Stars*, edited by L. Rezzolla, P. Pizzochero, D. I. Jones, N. Rea, and I. Vidaña (Springer International Publishing, 2018), pp. 337–400.
- [13] L. Tolos and L. Fabbietti, *Prog. Part. Nuc. Phys.* **112**, 103770 (2020).
- [14] J. J. Li, A. Sedrakian, and F. Weber, *Phys. Lett. B* **783**, 234–240 (2018).
- [15] W. M. Spinella and F. Weber, *Astron. Nachr.* **340**, 145–150 (2019).
- [16] G. Malfatti, M. G. Orsaria, I. F. Ranea-Sandoval, G. A. Contrera, and F. Weber, *Phys. Rev. D* **102**, 063008 (2020).
- [17] A. Sedrakian, J. J. Li, and F. Weber, *Prog. Part. Nucl. Phys.* **131**, 104041 (2023).
- [18] C. Alcock, E. Farhi, and A. Olinto, *Astrophys. J.* **310**, 261 (1986).
- [19] J. Madsen, in *Hadrons in Dense Matter and Hadrosynthesis*, edited by J. Cleymans, H. B. Geyer, and F. G. Scholtz (Springer Berlin Heidelberg, Berlin, Heidelberg, 1999), pp. 162–203.
- [20] F. Weber, *Prog. Part. Nucl. Phys.* **54**, 193–288 (2005).
- [21] M. Orsaria, H. Rodrigues, F. Weber, and G. A. Contrera, *Phys. Rev. C* **89**, 015806 (2014).
- [22] M. Alford, *Ann. Rev. Nucl. Part. Sci.* **51**, 131–160 (2001).
- [23] M. G. Alford, A. Schmitt, K. Rajagopal, and T. Schäfer, *Rev. Mod. Phys.* **80**, 1455–1515 (2008).
- [24] J. L. Zdunik and P. Haensel, *Astron. Astrophys.* **551**, A61 (2013).
- [25] M. Buballa, *Phys. Rep.* **407**, 205–376 (2005).
- [26] L. McLerran and S. Reddy, *Phys. Rev. Lett.* **122**, 122701 (2019).

- [27] G. Baym, *Phys. Rev. Lett.* **30**, 1340–1342 (1973).
- [28] D. Kaplan and A. Nelson, *Phys. Lett. B* **175**, 57–63 (1986).
- [29] N. K. Glendenning and J. Schaffner-Bielich, *Phys. Rev. C* **60**, 025803 (1999).
- [30] P. J. Ellis, R. Knorren, and M. Prakash, *Phys. Lett. B* **349**, 11–17 (1995).
- [31] A. Ramos, J. Schaffner-Bielich, and J. Wambach, in *Physics of Neutron Star Interiors* (Springer Berlin Heidelberg, 2001), pp. 175–202.
- [32] P. B. Demorest, T. Pennucci, S. M. Ransom, M. S. E. Roberts, and J. W. T. Hessels, *Nature* **467**, 1081–1083 (2010).
- [33] E. Fonseca *et al.*, *Astrophys. J.* **832**, 167 (2016).
- [34] Z. Arzoumanian *et al.*, *Astrophys. J. Suppl.* **235**, 37 (2018).
- [35] J. Antoniadis *et al.*, *Science* **340**, 6131 (2013).
- [36] H. T. Cromartie *et al.*, *Nat. Astron.* **4**, 72–76 (2020).
- [37] E. Fonseca *et al.*, *Astrophys. J. Lett.* **915**, L12 (2021).
- [38] B. P. Abbott *et al.*, *Phys. Rev. X* **9**, 011001 (2019).
- [39] B. P. Abbott *et al.*, *Astrophys. J. Lett.* **892**, L3 (2020).
- [40] T. E. Riley *et al.*, *Astrophys. J. Lett.* **887**, L21 (2019).
- [41] T. E. Riley *et al.*, *Astrophys. J. Lett.* **918**, L27 (2021).
- [42] M. C. Miller *et al.*, *Astrophys. J. Lett.* **887**, L24 (2019).
- [43] M. C. Miller *et al.*, *Astrophys. J. Lett.* **918**, L28 (2021).
- [44] R. W. Romani, D. Kandel, A. V. Filippenko, T. G. Brink, and W. Zheng, *Astrophys. J. Lett.* **934**, L17 (2022).
- [45] V. Doroshenko, V. Suleimanov, G. Pühlhofer, and A. Santangelo, *Nat. Astron.* **6**, 1444–1451 (2022).
- [46] D. Choudhury *et al.*, *Astrophys. J.* **submitted** (2024).
- [47] K. Fukushima and T. Hatsuda, *Rep. Prog. Phys.* **74**, 014001 (2011).
- [48] M. Alford, K. Rajagopal, and F. Wilczek, *Nucl. Phys. B* **537**, 443–458 (1999).
- [49] T. Gorda, A. Kurkela, R. Paatelainen, S. Säppi, and A. Vuorinen, *Phys. Rev. Lett.* **127**, 162003 (2021).
- [50] J. W. Holt, M. Rho, and W. Weise, *Phys. Rep.* **621**, 2–75 (2016).
- [51] C. Drischler, J. W. Holt, and C. Wellenhofer, *Annu. Rev. Nucl. Part. Sci.* **71**, 403–432 (2021).
- [52] C. Wellenhofer, J. W. Holt, and N. Kaiser, *Phys. Rev. C* **92**, 015801 (2015).
- [53] G. Aarts, *J. Phys. Conf. Ser.* **706**, 022004 (2016).
- [54] A. Lovato *et al.*, *arXiv:2211.02224* (2022).
- [55] L. McLerran and R. D. Pisarski, *Nucl. Phys. A* **796**, 83–100 (2007).
- [56] D. Logoteta, *Universe* **7**, 408 (2021).
- [57] T. Kojo, *AAPPS Bull.* **31**, 11 (2021).
- [58] G. Aarts *et al.*, *Prog. Part. Nucl. Phys.* **133**, 104070 (2023).

- [59] R. Kumar *et al.*, *Living Rev. Relativ.* **27**, 3 (2024).
- [60] C. Drischler, S. Han, J. M. Lattimer, M. Prakash, S. Reddy, and T. Zhao, *Phys. Rev. C* **103**, 045808 (2021).
- [61] C. Drischler, S. Han, and S. Reddy, *Phys. Rev. C* **105**, 035808 (2022).
- [62] O. Komoltsev and A. Kurkela, *Phys. Rev. Lett.* **128**, 202701 (2022).
- [63] A. W. Steiner, J. M. Lattimer, and E. F. Brown, *Astrophys. J.* **722**, 33–54 (2010).
- [64] F. Özel, D. Psaltis, T. Güver, G. Baym, C. Heinke, and S. Guillot, *Astrophys. J.* **820**, 28 (2016).
- [65] S. Bogdanov, C. O. Heinke, F. Özel, and T. Güver, *Astrophys. J.* **831**, 184 (2016).
- [66] C. A. Raithel, F. Özel, and D. Psaltis, *Astrophys. J. Lett.* **857**, L23 (2018).
- [67] S. K. Greif, G. Raaijmakers, K. Hebeler, A. Schwenk, and A. L. Watts, *Mon. Not. Roy. Astron. Soc.* **485**, 5363–5376 (2019).
- [68] G. Raaijmakers *et al.*, *Astrophys. J.* **887**, L22 (2019).
- [69] G. Raaijmakers *et al.*, *Astrophys. J.* **893**, L21 (2020).
- [70] G. Raaijmakers, S. K. Greif, K. Hebeler, T. Hinderer, S. Nisanke, A. Schwenk, T. E. Riley, A. L. Watts, J. M. Lattimer, and W. C. G. Ho, *Astrophys. J. Lett.* **918**, L29 (2021).
- [71] P. Landry and R. Essick, *Phys. Rev. D* **99**, 084049 (2019).
- [72] P. Landry, R. Essick, and K. Chatziioannou, *Phys. Rev. D* **101**, 123007 (2020).
- [73] R. Essick, I. Tews, P. Landry, S. Reddy, and D. E. Holz, *Phys. Rev. C* **102**, 055803 (2020).
- [74] I. Legred, K. Chatziioannou, R. Essick, S. Han, and P. Landry, *Phys. Rev. D* **104**, 063003 (2021).
- [75] C. A. Raithel, F. Özel, and D. Psaltis, *Astrophys. J.* **831**, 44 (2016).
- [76] C. A. Raithel, F. Özel, and D. Psaltis, *Astrophys. J.* **844**, 156 (2017).
- [77] C. A. Raithel, F. Özel, and D. Psaltis, *Astrophys. J.* **908**, 103 (2021).
- [78] T. Dietrich, M. W. Coughlin, P. T. H. Pang, M. Bulla, J. Heinzel, L. Issa, I. Tews, and S. Antier, *Science* **370**, 1450–1453 (2020).
- [79] M.-Z. Han, J.-L. Jiang, S.-P. Tang, and Y.-Z. Fan, *Astrophys. J.* **919**, 11 (2021).
- [80] P. T. H. Pang, I. Tews, M. W. Coughlin, M. Bulla, C. V. D. Broeck, and T. Dietrich, *Astrophys. J.* **922**, 14 (2021).
- [81] Y. Lim and J. W. Holt, *Eur. Phys. J. A* **55**, 209 (2019).
- [82] Y. Lim and J. W. Holt, *Galaxies* **10**, 99 (2022).
- [83] B. Biswas, *Astrophys. J.* **921**, 63 (2021).
- [84] B. Biswas, P. Char, R. Nandi, and S. Bose, *Phys. Rev. D* **103**, 103015 (2021).
- [85] B. Biswas and S. Datta, *Phys. Rev. D* **106**, 043012 (2022).
- [86] T. Gorda, O. Komoltsev, and A. Kurkela, *Astrophys. J.* **950**, 107 (2023).
- [87] M. Al-Mamun, A. W. Steiner, J. Nättilä, J. Lange, R. O’Shaughnessy, I. Tews, S. Gandolfi, C. Heinke, and S. Han, *Phys. Rev. Lett.* **126**, 061101 (2021).

- [88] S. Huth *et al.*, *Nature* **606**, 276–280 (2022).
- [89] I. Legred, K. Chatziioannou, R. Essick, and P. Landry, *Phys. Rev. D* **105**, 043016 (2022).
- [90] M. Marczenko, L. McLerran, K. Redlich, and C. Sasaki, *Phys. Rev. C* **107**, 025802 (2023).
- [91] E. Annala, T. Gorda, J. Hirvonen, O. Komoltsev, A. Kurkela, J. Nättilä, and A. Vuorinen, *Nat. Commun.* **14**, 8451 (2023).
- [92] M.-Z. Han, Y.-J. Huang, S.-P. Tang, and Y.-Z. Fan, *Sci. Bull.* **68**, 913–919 (2023).
- [93] R. Essick, I. Legred, K. Chatziioannou, S. Han, and P. Landry, *Phys. Rev. D* **108**, 043013 (2023).
- [94] J. Takátsy, P. Kovács, G. Wolf, and J. Schaffner-Bielich, *Phys. Rev. D* **108**, 043002 (2023).
- [95] J.-L. Jiang, C. Ecker, and L. Rezzolla, *Astrophys. J.* **949**, 11 (2023).
- [96] L. Brandes, W. Weise, and N. Kaiser, *Phys. Rev. D* **107**, 014011 (2023).
- [97] L. Brandes, W. Weise, and N. Kaiser, *Phys. Rev. D* **108**, 094014 (2023).
- [98] D. Mroczek, M. C. Miller, J. Noronha-Hostler, and N. Yunes, *arXiv:2309.02345* (2023).
- [99] P. T. H. Pang, L. Sivertsen, R. Somasundaram, T. Dietrich, S. Sen, I. Tews, M. W. Coughlin, and C. Van Den Broeck, *Phys. Rev. C* **109**, 025807 (2024).
- [100] Y. Lim and A. Schwenk, *Phys. Rev. C* **109**, 035801 (2024).
- [101] C. Y. Tsang, M. B. Tsang, W. G. Lynch, R. Kumar, and C. J. Horowitz, *Nat. Astron.* **8**, 328–336 (2024).
- [102] Y.-Z. Fan, M.-Z. Han, J.-L. Jiang, D.-S. Shao, and S.-P. Tang, *Phys. Rev. D* **109**, 043052 (2024).
- [103] H. Koehn *et al.*, *arXiv:2402.04172* (2024).
- [104] D. D. Ofengeim, P. S. Shternin, and T. Piran, *arXiv:2404.17647* (2024).
- [105] O. Komoltsev, *arXiv:2404.05637* (2024).
- [106] T. Malik, V. Dexheimer, and C. Providência, *arXiv:2404.07936* (2024).
- [107] S.-P. Tang, Y.-J. Huang, M.-Z. Han, and Y.-Z. Fan, *arXiv:2404.09563* (2024).
- [108] Y. Fujimoto, K. Fukushima, and K. Murase, *Phys. Rev. D* **98**, 023019 (2018).
- [109] Y. Fujimoto, K. Fukushima, and K. Murase, *Phys. Rev. D* **101**, 054016 (2020).
- [110] Y. Fujimoto, K. Fukushima, and K. Murase, *J. High Energy Phys.* **2021**, 273 (2021).
- [111] M. Ferreira and C. Providência, *JCAP* **07**, 011 (2021).
- [112] F. Morawski and M. Bejger, *Astron. Astrophys.* **642**, A78 (2020).
- [113] S. Traversi and P. Char, *Astrophys. J.* **905**, 9 (2020).
- [114] P. G. Krastev, *Galaxies* **10**, 16 (2022).
- [115] F. Morawski and M. Bejger, *Phys. Rev. C* **106**, 065802 (2022).
- [116] M. Ferreira, V. Carvalho, and C. Providência, *Phys. Rev. D* **106**, 103023 (2022).

- [117] A. Thete, K. Banerjee, and T. Malik, *Phys. Rev. D* **108**, 063028 (2023).
- [118] S. Soma, L. Wang, S. Shi, H. Stöcker, and K. Zhou, *JCAP* **08**, 071 (2022).
- [119] S. Soma, L. Wang, S. Shi, H. Stöcker, and K. Zhou, *Phys. Rev. D* **107**, 083028 (2023).
- [120] P. G. Krastev, *Symmetry* **15**, 1123 (2023).
- [121] L.-J. Guo, J.-Y. Xiong, Y. Ma, and Y.-L. Ma, *Astrophys. J.* **965**, 47 (2024).
- [122] S. Chatterjee, H. Sudhakaran, and R. Mallick, *arXiv:2302.13648* (2023).
- [123] W. Zhou, J. Hu, Y. Zhang, and H. Shen, *Astrophys. J.* **950**, 186 (2023).
- [124] V. Carvalho, M. Ferreira, T. Malik, and C. Providência, *Phys. Rev. D* **108**, 043031 (2023).
- [125] V. Carvalho, M. Ferreira, and C. Providência, *Phys. Rev. D* **109**, 123038 (2024).
- [126] Y. Fujimoto, K. Fukushima, S. Kamata, and K. Murase, *arXiv:2401.12688* (2024).
- [127] D. Farrell, P. Baldi, J. Ott, A. Ghosh, A. W. Steiner, A. Kavitkar, L. Lindblom, D. Whiteson, and F. Weber, *JCAP* **02**, 016 (2023).
- [128] D. Farrell, P. Baldi, J. Ott, A. Ghosh, A. W. Steiner, A. Kavitkar, L. Lindblom, D. Whiteson, and F. Weber, *JCAP* **12**, 022 (2023).
- [129] K. Hebel, J. M. Lattimer, C. J. Pethick, and A. Schwenk, *Astrophys. J.* **773**, 11 (2013).
- [130] E. Annala, T. Gorda, A. Kurkela, J. Nättilä, and A. Vuorinen, *Nat. Phys.* **16**, 907–910 (2020).
- [131] E. Annala, T. Gorda, E. Katerini, A. Kurkela, J. Nättilä, V. Paschalidis, and A. Vuorinen, *Phys. Rev. X* **12**, 011058 (2022).
- [132] S. Altiparmak, C. Ecker, and L. Rezzolla, *Astrophys. J. Lett.* **939**, L34 (2022).
- [133] C. Ecker and L. Rezzolla, *Astrophys. J. Lett.* **939**, L35 (2022).
- [134] C. Ecker and L. Rezzolla, *Mon. Not. Roy. Astron. Soc.* **519**, 2615–2622 (2022).
- [135] T. Gorda, K. Hebel, A. Kurkela, A. Schwenk, and A. Vuorinen, *Astrophys. J.* **955**, 100 (2023).
- [136] R. Somasundaram, I. Tews, and J. Margueron, *Phys. Rev. C* **107**, 025801 (2023).
- [137] R. Somasundaram, I. Tews, and J. Margueron, *Phys. Rev. C* **107**, 1052801 (2023).
- [138] I. Tews, J. Margueron, and S. Reddy, *Eur. Phys. J. A* **55**, 97 (2019).
- [139] T. E. Riley, G. Raaijmakers, and A. L. Watts, *Mon. Not. Roy. Astron. Soc.* **478**, 1093–1131 (2018).
- [140] M. C. Miller, C. Chirenti, and F. K. Lamb, *Astrophys. J.* **888**, 12 (2019).
- [141] A. W. Steiner, C. O. Heinke, S. Bogdanov, C. Li, W. C. G. Ho, A. Bahramian, and S. Han, *Mon. Not. Roy. Astron. Soc.* **476**, 421–435 (2018).
- [142] A. Boehnlein *et al.*, *Rev. Mod. Phys.* **94**, 031003 (2022).
- [143] K. Zhou, L. Wang, L.-G. Pang, and S. Shi, *Prog. Part. Nucl. Phys.* **135**, 104084 (2024).
- [144] G. Papamakarios, D. Sterratt, and I. Murray, *arXiv:1805.07226* (2019).
- [145] L. Dinh, J. Sohl-Dickstein, and S. Bengio, *arXiv:1605.08803* (2016).

- [146] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan, *J. Mach. Learn. Res.* **22** (2021).
- [147] R. M. Neal, in *Handbook of Markov Chain Monte Carlo* (Chapman and Hall/CRC, May 2011) Chap. 5.
- [148] M. Betancourt, [arXiv:1701.02434](https://arxiv.org/abs/1701.02434) (2017).
- [149] A. Burrows and J. M. Lattimer, *Astrophys. J.* **307**, 178 (1986).
- [150] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, and I. Sutskever, [arXiv:2102.12092](https://arxiv.org/abs/2102.12092) (2021).
- [151] S. Chandrasekhar, *Rev. Mod. Phys.* **56**, 137–147 (1984).
- [152] J. W. T. Hessels, S. M. Ransom, I. H. Stairs, P. C. C. Freire, V. M. Kaspi, and F. Camilo, *Science* **311**, 1901–1904 (2006).
- [153] A. P. Igoshev, S. B. Popov, and R. Hollerbach, *Universe* **7**, 351 (2021).
- [154] D. R. Lorimer, *Living Rev. Relativ.* **8**, 7 (2005).
- [155] R. N. Manchester, G. B. Hobbs, A. Teoh, and M. Hobbs, *Astron. J.* **129**, 1993–2006 (2005).
- [156] Y. Suwa, T. Yoshida, M. Shibata, H. Umeda, and K. Takahashi, *Mon. Not. Roy. Astron. Soc.* **481**, 3305–3312 (2018).
- [157] A. Sorensen *et al.*, *Prog. Part. Nucl. Phys.* **134**, 104080 (2024).
- [158] J. M. Lattimer, *AIP Conf. Proc.* **1645**, 61–78 (2015).
- [159] D. Yakovlev and C. Pethick, *Ann. Rev. Astron. Astr.* **42**, 169–210 (2004).
- [160] S. Ascenzi, V. Graber, and N. Rea, *Astropart. Phys.* **158**, 102935 (2024).
- [161] J. M. Lattimer and M. Prakash, *Science* **304**, 536–542 (2004).
- [162] A. Einstein, *Annalen der Physik* **354**, 769–822 (1916).
- [163] L. Brandes and W. Weise, *Symmetry* **16**, 111 (2024).
- [164] N. K. Glendenning, *Compact stars, nuclear physics, particle physics, and general relativity* (Springer, 1997), p. 390.
- [165] S. M. Carroll, *Spacetime and Geometry, An Introduction to General Relativity* (Pearson Education Limited, 2014), p. 528.
- [166] L. Lindblom, *Astrophys. J.* **398**, 569 (1992).
- [167] P. Haensel, J. L. Zdunik, M. Bejger, and J. M. Lattimer, *Astron. Astrophys.* **502**, 605–610 (2009).
- [168] E. E. Flanagan and T. Hinderer, *Phys. Rev. D* **77**, 021502 (2008).
- [169] T. Hinderer, *Astrophys. J.* **677**, 1216–1220 (2008).
- [170] K. Chatzioannou, *Gen. Rel. Grav.* **52**, 109 (2020).
- [171] T. Hinderer, B. D. Lackey, R. N. Lang, and J. S. Read, *Phys. Rev. D* **81**, 123016 (2010).
- [172] LIGO Scientific Collaboration, *LIGO Algorithm Library - LALSuite*, free software (GPL), 2018.

- [173] M. E. Peskin and D. V. Schroeder, *An introduction to quantum field theory* (Addison-Wesley Pub. Co., 1995).
- [174] R. L. Workman *et al.*, *Prog. Theor. Exp. Phys.* **2022**, 083C01 (2022).
- [175] C. Gattringer and C. B. Lang, *Quantum Chromodynamics on the Lattice: An Introductory Presentation* (Springer Berlin Heidelberg, 2010).
- [176] C. S. Fischer, *Prog. Part. Nuc. Phys.* **105**, 1–60 (2019).
- [177] N. Dupuis, L. Canet, A. Eichhorn, W. Metzner, J. M. Pawłowski, M. Tissier, and N. Wschebor, *Phys. Rep.* **910**, 1–114 (2021).
- [178] E. Epelbaum, H.-W. Hammer, and U.-G. Meißner, *Rev. Mod. Phys.* **81**, 1773–1825 (2009).
- [179] E. Epelbaum, [arXiv:1001.3229](https://arxiv.org/abs/1001.3229) (2010).
- [180] S. Weinberg, *Physica A* **96**, 327–340 (1979).
- [181] E. Jenkins and A. V. Manohar, *Phys. Lett. B* **255**, 558–562 (1991).
- [182] V. Bernard, N. Kaiser, and U.-G. Meißner, *Int. J. Mod. Phys. E* **04**, 193–344 (1995).
- [183] S. Gandolfi, J. Carlson, and S. Reddy, *Phys. Rev. C* **85**, 032801 (2012).
- [184] S. Gandolfi, J. Carlson, S. Reddy, A. W. Steiner, and R. B. Wiringa, *Eur. Phys. J. A* **50**, 10 (2014).
- [185] J. B. Elliott, P. T. Lake, L. G. Moretto, and L. Phair, *Phys. Rev. C* **87**, 054622 (2013).
- [186] J. A. Melendez, R. J. Furnstahl, D. R. Phillips, M. T. Pratola, and S. Wesolowski, *Phys. Rev. C* **100**, 044001 (2019).
- [187] C. Drischler, R. Furnstahl, J. Melendez, and D. Phillips, *Phys. Rev. Lett.* **125**, 202702 (2020).
- [188] C. Drischler, J. A. Melendez, R. J. Furnstahl, and D. R. Phillips, *Phys. Rev. C* **102**, 054315 (2020).
- [189] J. Keller, K. Hebeler, and A. Schwenk, *Phys. Rev. Lett.* **130**, 072701 (2023).
- [190] I. Tews, J. Carlson, S. Gandolfi, and S. Reddy, *Astrophys. J.* **860**, 149 (2018).
- [191] J. I. Kapusta and C. Gale, *Finite-Temperature Field Theory: Principles and Applications* (Cambridge University Press, Aug. 2006).
- [192] M. Laine and A. Vuorinen, *Basics of Thermal Field Theory: A Tutorial on Perturbative Computations* (Springer International Publishing, 2016).
- [193] J. Ghiglieri, A. Kurkela, M. Strickland, and A. Vuorinen, *Phys. Rep.* **880**, 1–73 (2020).
- [194] B. A. Freedman and L. D. McLerran, *Phys. Rev. D* **16**, 1147–1168 (1977).
- [195] B. A. Freedman and L. D. McLerran, *Phys. Rev. D* **16**, 1169–1185 (1977).
- [196] A. Kurkela, P. Romatschke, and A. Vuorinen, *Phys. Rev. D* **81**, 105021 (2010).
- [197] E. S. Fraga, A. Kurkela, and A. Vuorinen, *Astrophys. J. Lett.* **781**, L25 (2014).

- [198] T. Gorda, A. Kurkela, R. Paatelainen, S. Säppi, and A. Vuorinen, *Phys. Rev. D* **104**, 074015 (2021).
- [199] T. Gorda, R. Paatelainen, S. Säppi, and K. Seppänen, *Phys. Rev. Lett.* **131**, 181902 (2023).
- [200] M. Hippert, E. Fraga, and J. Noronha, *Phys. Rev. D* **104**, 034011 (2021).
- [201] P. Bedaque and A. W. Steiner, *Phys. Rev. Lett.* **114**, 031103 (2015).
- [202] M. Leonhardt, M. Pospiech, B. Schallmo, J. Braun, C. Drischler, K. Hebeler, and A. Schwenk, *Phys. Rev. Lett.* **125**, 142502 (2020).
- [203] K. Iida and E. Itou, *Prog. Theor. Exp. Phys.* **2022**, 111B01 (2022).
- [204] B. B. Brandt, F. Cuteri, and G. Endrődi, *J. High Energy Phys.* **2023**, 55 (2023).
- [205] R. Abbott, W. Detmold, F. Romero-López, Z. Davoudi, M. Illa, A. Parreño, R. J. Perry, P. E. Shanahan, and M. L. Wagman, *Phys. Rev. D* **108**, 114506 (2023).
- [206] Y. Fujimoto, K. Fukushima, L. D. McLerran, and M. Praszalowicz, *Phys. Rev. Lett.* **129**, 252702 (2022).
- [207] Y. Fujimoto and K. Fukushima, *Phys. Rev. D* **105**, 014025 (2022).
- [208] K. Fukushima and C. Sasaki, *Prog. Part. Nucl. Phys.* **72**, 99–154 (2013).
- [209] A. Bazavov *et al.*, *Phys. Rev. D* **90**, 094503 (2014).
- [210] S. Borsanyi, Z. Fodor, C. Hoelbling, S. D. Katz, S. Krieg, and K. K. Szabo, *Phys. Lett. B* **730**, 99–104 (2014).
- [211] P. Braun-Munzinger, V. Koch, T. Schäfer, and J. Stachel, *Phys. Rep.* **621**, 76–126 (2016).
- [212] A. Andronic, P. Braun-Munzinger, K. Redlich, and J. Stachel, *Nature* **561**, 321–330 (2018).
- [213] P. Romatschke and U. Romatschke, *Relativistic Fluid Dynamics in and out of Equilibrium* (Cambridge University Press, 2019).
- [214] T. Schäfer and F. Wilczek, *Phys. Rev. D* **60**, 114033 (1999).
- [215] A. Geißel, T. Gorda, and J. Braun, [arXiv:2403.18010](https://arxiv.org/abs/2403.18010) (2024).
- [216] F. J. Llanes-Estrada and E. Lope-Oter, *Prog. Part. Nucl. Phys.* **109**, 103715 (2019).
- [217] S. Benić, D. Blaschke, D. E. Alvarez-Castillo, T. Fischer, and S. Typel, *Astron. Astrophys.* **577**, A40 (2015).
- [218] S. Han, M. Mamun, S. Lalit, C. Constantinou, and M. Prakash, *Phys. Rev. D* **100**, 103022 (2019).
- [219] M. Drews and W. Weise, *Phys. Rev. C* **91**, 035802 (2015).
- [220] M. Drews and W. Weise, *Prog. Part. Nucl. Phys.* **93**, 69–107 (2017).
- [221] L. Brandes, N. Kaiser, and W. Weise, *Eur. Phys. J. A* **57**, 243 (2021).
- [222] I. F. Ranea-Sandoval, S. Han, M. G. Orsaria, G. A. Contrera, F. Weber, and M. G. Alford, *Phys. Rev. C* **93**, 045812 (2016).
- [223] A. Bzdak, S. Esumi, V. Koch, J. Liao, M. Stephanov, and N. Xu, *Phys. Rep.* **853**, 1–87 (2020).

- [224] A. Bazavov *et al.*, *Phys. Lett. B* **795**, 15–21 (2019).
- [225] S. Borsanyi, Z. Fodor, J. N. Guenther, R. Kara, S. D. Katz, P. Parotto, A. Pasztor, C. Ratti, and K. K. Szabó, *Phys. Rev. Lett.* **125**, 052001 (2020).
- [226] W.-j. Fu, J. M. Pawłowski, and F. Rennecke, *Phys. Rev. D* **101**, 054032 (2020).
- [227] F. Gao and J. M. Pawłowski, *Phys. Lett. B* **820**, 136584 (2021).
- [228] P. J. Gunkel and C. S. Fischer, *Phys. Rev. D* **104**, 054022 (2021).
- [229] W.-j. Fu, *Commun. Theor. Phys* **74**, 097304 (2022).
- [230] R. Critelli, J. Noronha, J. Noronha-Hostler, I. Portillo, C. Ratti, and R. Rougemont, *Phys. Rev. D* **96**, 096026 (2017).
- [231] M. Hippert, J. Grefa, T. A. Manning, J. Noronha, J. Noronha-Hostler, I. P. Vazquez, C. Ratti, R. Rougemont, and M. Trujillo, [arXiv:2309.00579](https://arxiv.org/abs/2309.00579) (2023).
- [232] G. Baym, S. Furusawa, T. Hatsuda, T. Kojo, and H. Togashi, *Astrophys. J.* **885**, 42 (2019).
- [233] K. Fukushima, T. Kojo, and W. Weise, *Phys. Rev. D* **102**, 096017 (2020).
- [234] T. Kojo, *Phys. Rev. D* **104**, 074005 (2021).
- [235] T. Kojo, G. Baym, and T. Hatsuda, *Astrophys. J.* **934**, 46 (2022).
- [236] T. Zhao and J. M. Lattimer, *Phys. Rev. D* **102**, 023021 (2020).
- [237] Y. Fujimoto, T. Kojo, and L. D. McLerran, *Phys. Rev. Lett.* **132**, 112701 (2024).
- [238] D. Lonardonì, A. Lovato, S. Gandolfi, and F. Pederiva, *Phys. Rev. Lett.* **114**, 092301 (2015).
- [239] J. Haidenbauer, U. G. Meißner, N. Kaiser, and W. Weise, *Eur. Phys. J. A* **53**, 121 (2017).
- [240] D. Gerstung, N. Kaiser, and W. Weise, *Eur. Phys. J. A* **56**, 175 (2020).
- [241] E. Friedman and A. Gal, *Phys. Lett. B* **837**, 137669 (2023).
- [242] T. Motta, P. Guichon, and A. Thomas, *Nucl. Phys. A* **1009**, 122157 (2021).
- [243] L. Brandes, C. Modi, A. Ghosh, D. Farrell, L. Lindblom, L. Heinrich, A. W. Steiner, F. Weber, and D. Whiteson, *JCAP* **09**, 009 (2024).
- [244] I. Goodfellow, Y. Bengio, A. Courville, and F. Bach, *Deep Learning* (MIT Press, 2017).
- [245] A. Zhang, Z. C. Lipton, M. Li, and A. J. Smola, *Dive into Deep Learning* (Cambridge University Press, 2023).
- [246] K. Hornik, M. Stinchcombe, and H. White, *Neural Netw.* **2**, 359–366 (1989).
- [247] D. P. Kingma and J. Ba, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014).
- [248] G. Papamakarios, T. Pavlakou, and I. Murray, *NeurIPS* **30** (2017).
- [249] D. P. Kingma, T. Salimans, R. Jozefowicz, X. Chen, I. Sutskever, and M. Welling, *NeurIPS* **29** (2016).
- [250] M. Germain, K. Gregor, I. Murray, and H. Larochelle, *ICML* **37**, 881–889 (2015).
- [251] S. R. Green, C. Simpson, and J. Gair, *Phys. Rev. D* **102**, 104057 (2020).
- [252] F. Karsch, A. Lahiri, M. Neumann, and C. Schmidt, *PoS LATTICE2022*, 027 (2022).

- [253] C. Krause and D. Shih, *Phys. Rev. D* **107**, 113004 (2023).
- [254] J. S. Read, B. D. Lackey, B. J. Owen, and J. L. Friedman, *Phys. Rev. D* **79**, 124032 (2009).
- [255] L. Lindblom, *Phys. Rev. D* **82**, 103011 (2010).
- [256] H. Tan, T. Dore, V. Dexheimer, J. Noronha-Hostler, and N. Yunes, *Phys. Rev. D* **105**, 023018 (2022).
- [257] G. Baym, C. Pethick, and P. Sutherland, *Astrophys. J.* **170**, 299 (1971).
- [258] R. Gamba, J. S. Read, and L. E. Wade, *Class. Quantum Grav.* **37**, 025008 (2019).
- [259] G. Grams, R. Somasundaram, J. Margueron, and S. Reddy, *Phys. Rev. C* **105**, 035806 (2022).
- [260] L. Lindblom, *Phys. Rev. D* **97**, 123019 (2018).
- [261] O. Komoltsev, R. Somasundaram, T. Gorda, A. Kurkela, J. Margueron, and I. Tews, *Phys. Rev. D* **109**, 094030 (2024).
- [262] B. W. Silverman, *Density Estimation for Statistics and Data Analysis* (Chapman and Hall, London, 1986).
- [263] H. Jeffreys, *Theory of Probability* (Oxford University Press, New York, Aug. 1961), 472 pp.
- [264] M. D. Lee and E.-J. Wagenmakers, *Bayesian Cognitive Modeling* (Cambridge University Press, Cambridge, England, Apr. 2014), 280 pp.
- [265] I. Mandel, W. M. Farr, and J. R. Gair, *Mon. Not. Roy. Astron. Soc.* **486**, 1086–1093 (2019).
- [266] D. Wysocki, R. O’Shaughnessy, L. Wade, and J. Lange, [arXiv:2001.01747](https://arxiv.org/abs/2001.01747) (2020).
- [267] S. Typel, G. Röpke, T. Klähn, D. Blaschke, and H. H. Wolter, *Phys. Rev. C* **81**, 015803 (2010).
- [268] A. Akmal, V. R. Pandharipande, and D. G. Ravenhall, *Phys. Rev. C* **58**, 1804–1828 (1998).
- [269] M. C. Miller, in *Timing neutron stars: pulsations, oscillations and explosions*, Vol. 461, Astrophysics and Space Science Library (Springer, 2020), pp. 1–51.
- [270] P. C. C. Freire, [arXiv:0907.3219](https://arxiv.org/abs/0907.3219) (2009).
- [271] M. Coleman Miller and F. K. Lamb, *Eur. Phys. J. A* **52**, 63 (2016).
- [272] A. Konstantinou and S. M. Morsink, *Astrophys. J.* **934**, 139 (2022).
- [273] A. G. Sullivan and R. W. Romani, [arXiv:2405.13889](https://arxiv.org/abs/2405.13889) (2024).
- [274] M. Linares, T. Shahbaz, and J. Casares, *Astrophys. J.* **859**, 54 (2018).
- [275] D. Kandel and R. W. Romani, *Astrophys. J.* **942**, 6 (2022).
- [276] A. L. Watts, *AIP Conf. Proc.* **2127**, 020008 (2019).
- [277] A. L. Watts *et al.*, *Rev. Mod. Phys.* **88**, 021001 (2016).
- [278] F. Özel, D. Psaltis, Z. Arzoumanian, S. Morsink, and M. Bauböck, *Astrophys. J.* **832**, 92 (2016).
- [279] S. Vinciguerra *et al.*, *Astrophys. J.* **961**, 62 (2024).
- [280] T. Salmi *et al.*, *Astrophys. J.* **941**, 150 (2022).

- [281] A. J. Dittmann *et al.*, [arXiv:2406.14467](#) (2024).
- [282] D. J. Reardon *et al.*, *Mon. Not. Roy. Astron. Soc.* **455**, 1751–1769 (2015).
- [283] S. Guillot *et al.*, *Astrophys. J. Lett.* **887**, L27 (2019).
- [284] A. L. Watts *et al.*, *Sci. China Phys. Mech. Astron.* **62**, 29503 (2018).
- [285] P. S. Ray *et al.*, [arXiv:1903.03035](#) (2019).
- [286] C. O. Heinke, J. E. Grindlay, P. M. Lugger, H. N. Cohn, P. D. Edmonds, D. A. Lloyd, and A. M. Cool, *Astrophys. J.* **598**, 501–515 (2003).
- [287] K. G. Elshamouty, C. O. Heinke, S. M. Morsink, S. Bogdanov, and A. L. Stevens, *Astrophys. J.* **826**, 162 (2016).
- [288] D. K. Galloway and L. Keek, in *Astrophysics and Space Science Library* (Springer Berlin Heidelberg, Oct. 2020), pp. 209–262.
- [289] J. Nättilä, A. W. Steiner, J. J. E. Kajava, V. F. Suleimanov, and J. Poutanen, *Astron. Astrophys.* **591**, A25 (2016).
- [290] J. M. Lattimer, *Universe* **5**, 159 (2019).
- [291] J. Nättilä, M. C. Miller, A. W. Steiner, J. J. E. Kajava, V. F. Suleimanov, and J. Poutanen, *Astron. Astrophys.* **608**, A31 (2017).
- [292] T. M. Tauris and H.-T. Janka, *Astrophys. J.* **886**, L20 (2019).
- [293] J. A. J. Alford and J. P. Halpern, *Astrophys. J.* **944**, 36 (2023).
- [294] T. Dietrich, T. Hinderer, and A. Samajdar, *Gen. Rel. Grav.* **53**, 27 (2021).
- [295] B. P. Abbott *et al.*, *Phys. Rev. Lett.* **121**, 161101 (2018).
- [296] S. De, D. Finstad, J. M. Lattimer, D. A. Brown, E. Berger, and C. M. Biwer, *Phys. Rev. Lett.* **121**, 091102 (2018).
- [297] C. D. Capano, I. Tews, S. M. Brown, B. Margalit, S. De, S. Kumar, D. A. Brown, B. Krishnan, and S. Reddy, *Nat. Astron.* **4**, 625–632 (2020).
- [298] M. Fasano, T. Abdelsalhin, A. Maselli, and V. Ferrari, *Phys. Rev. Lett.* **123**, 141101 (2019).
- [299] B. P. Abbott *et al.*, *Astrophys. J. Lett.* **848**, L12 (2017).
- [300] G. Ghirlanda *et al.*, *Science* **363**, 968–971 (2019).
- [301] B. D. Metzger, *Living Rev. Relativ.* **23**, 1 (2019).
- [302] L. Rezzolla, E. R. Most, and L. R. Weih, *Astrophys. J.* **852**, L25 (2018).
- [303] B. Margalit and B. D. Metzger, *Astrophys. J.* **850**, L19 (2017).
- [304] M. Ruiz, S. L. Shapiro, and A. Tsokaros, *Phys. Rev. D* **97**, 021501 (2018).
- [305] A. Colombo, O. S. Salafia, F. Gabrielli, G. Ghirlanda, B. Giacomazzo, and M. C. A. Perego, *Astrophys. J.* **937**, 79 (2022).
- [306] S. Hild *et al.*, *Class. Quantum Grav.* **28**, 094013 (2011).
- [307] E. D. Hall, *Galaxies* **10**, 90 (2022).
- [308] F. Iacovelli, M. Mancarella, C. Mondal, A. Puecher, T. Dietrich, F. Gulminelli, M. Maggiore, and M. Oertel, *Phys. Rev. D* **108**, 122006 (2023).
- [309] T. Zhao and J. M. Lattimer, *Phys. Rev. D* **106**, 123002 (2022).

- [310] B. K. Pradhan, D. Pathak, and D. Chatterjee, *Astrophys. J.* **956**, 38 (2023).
- [311] H. Sotani, N. Yasutake, T. Maruyama, and T. Tatsumi, *Phys. Rev. D* **83**, 024014 (2011).
- [312] M. Sieniawska and D. I. Jones, *Mon. Not. Roy. Astron. Soc.* **509**, 5179–5187 (2021).
- [313] P. Landry and B. Kumar, *Astrophys. J.* **868**, L22 (2018).
- [314] D. Adhikari *et al.*, *Phys. Rev. Lett.* **126**, 172502 (2021).
- [315] B. T. Reed, F. Fattoyev, C. Horowitz, and J. Piekarewicz, *Phys. Rev. Lett.* **126**, 172503 (2021).
- [316] R. Essick, P. Landry, A. Schwenk, and I. Tews, *Phys. Rev. C* **104**, 065804 (2021).
- [317] D. Adhikari *et al.*, *Phys. Rev. Lett.* **129**, 042501 (2022).
- [318] M. Salinas and J. Piekarewicz, *Phys. Rev. C* **107**, 045802 (2023).
- [319] J. Steinheimer, A. Motornenko, A. Sorensen, Y. Nara, V. Koch, and M. Bleicher, *Eur. Phys. J. C* **82**, 911 (2022).
- [320] A. L. Fèvre, Y. Leifels, W. Reisdorf, J. Aichelin, and C. Hartnack, *Nucl. Phys. A* **945**, 112–133 (2016).
- [321] P. Russotto *et al.*, *Phys. Rev. C* **94**, 034608 (2016).
- [322] C. O. Heinke, G. B. Rybicki, R. Narayan, and J. E. Grindlay, *Astrophys. J.* **644**, 1090–1103 (2006).
- [323] K. Arnaud, *ASP Conf. Ser.* 101 **17** (1996).
- [324] J. M. Lattimer and A. W. Steiner, *Astrophys. J.* **784**, 123 (2014).
- [325] K. Cranmer, J. Brehmer, and G. Louppe, *PNAS* **117**, 30055–30062 (2020).
- [326] J.-M. Lueckmann, J. Boelts, D. S. Greenberg, P. J. Gonçalves, and J. H. Macke, *arXiv:2101.04653* (2021).
- [327] A. Delaunoy, A. Wehenkel, T. Hinderer, S. Nissanke, C. Weniger, A. R. Williamson, and G. Louppe, *arXiv:2010.12931* (2020).
- [328] M. Dax, S. R. Green, J. Gair, J. H. Macke, A. Buonanno, and B. Schölkopf, *Phys. Rev. Lett.* **127**, 241103 (2021).
- [329] S. R. Green and J. Gair, *MLST* **2**, 03LT01 (2021).
- [330] K. Cranmer, J. Pavez, and G. Louppe, *arXiv:1506.02169* (2015).
- [331] J. Brehmer, F. Kling, I. Espejo, and K. Cranmer, *Comput. Softw. Big Sci.* **4**, 3 (2020).
- [332] A. Ghosh, *hal-02971995*, 171–176 (2020).
- [333] J. Alsing, T. Charnock, S. Feeney, and B. Wandelt, *Mon. Not. Roy. Astron. Soc.* **488**, 4440–4458 (2019).
- [334] N. Jeffrey, J. Alsing, and F. Lanusse, *Mon. Not. Roy. Astron. Soc.* **501**, 954–969 (2021).
- [335] C. Hahn, M. Eickenberg, S. Ho, J. Hou, P. Lemos, E. Massara, C. Modi, A. M. Dizgah, B. R.-S. Blancard, and M. M. Abidi, *PNAS* **120**, e2218810120 (2023).

- [336] C. Modi and O. H. E. Philcox, [arXiv:2309.10270 \(2023\)](#).
- [337] C. Modi, S. Pandey, M. Ho, C. Hahn, B. R.-S. Blancard, and B. Wandelt, [arXiv:2309.15071 \(2023\)](#).
- [338] A. Tejero-Cantero, J. Boelts, M. Deistler, J.-M. Lueckmann, C. Durkan, P. J. Gonçalves, D. S. Greenberg, and J. H. Macke, [Journal of Open Source Software](#) **5**, 2505 (2020).
- [339] J. Hermans, A. Delaunoy, F. Rozet, A. Wehenkel, and G. Louppe, [arXiv:2110.06581 \(2021\)](#).
- [340] G. Papamakarios and I. Murray, [arXiv:1605.06376 \(2018\)](#).
- [341] D. S. Greenberg, M. Nonnenmacher, and J. H. Macke, [arXiv:1905.07488 \(2019\)](#).
- [342] C. Durkan, I. Murray, and G. Papamakarios, [arXiv:2002.03712 \(2020\)](#).
- [343] J. Hermans, V. Begy, and G. Louppe, [arXiv:1903.04057 \(2020\)](#).
- [344] L. Brandes, K. Fukushima, K. Iida, and C. Yu, *in preparation* (2024).
- [345] S. Typel, M. Oertel, and T. Klähn, [Phys. Part. Nuclei](#) **46**, 633–664 (2015).
- [346] M. Oertel, M. Hempel, T. Klähn, and S. Typel, [Rev. Mod. Phys.](#) **89**, 015007 (2017).
- [347] Y. Fujimoto, [Phys. Rev. D](#) **109**, 054035 (2024).
- [348] S. Fiorilla, N. Kaiser, and W. Weise, [Nucl. Phys. A](#) **880**, 65–87 (2012).
- [349] C. Wellenhofer, J. W. Holt, N. Kaiser, and W. Weise, [Phys. Rev. C](#) **89**, 064009 (2014).
- [350] L. Tsaloukidis, P. Koliogiannis, A. Kanakis-Pegios, and C. Moustakidis, [Phys. Rev. D](#) **107**, 023012 (2023).
- [351] B. Friman and W. Weise, [Phys. Rev. C](#) **100**, 065807 (2019).
- [352] R. Abbott *et al.*, [Astrophys. J. Lett.](#) **896**, L44 (2020).
- [353] S. Cotter, P. Giudice, S. Hands, and J.-I. Skullerud, [Phys. Rev. D](#) **87**, 034507 (2013).
- [354] M. Alford, M. Braby, M. Paris, and S. Reddy, [Astrophys. J.](#) **629**, 969–978 (2005).
- [355] H. Satz, [Phys. Rept.](#) **88**, 349 (1982).
- [356] E. Farhi and R. L. Jaffe, [Phys. Rev. D](#) **30**, 2379–2390 (1984).
- [357] P. Guichon, [Phys. Lett. B](#) **200**, 235–240 (1988).
- [358] A. Aziz, S. Ray, F. Rahaman, M. Khlopov, and B. K. Guha, [Int. J. Mod. Phys. D](#) **28**, 1941006 (2019).
- [359] R. Chiba and T. Kojo, [Phys. Rev. D](#) **109**, 076006 (2024).
- [360] E. Oset, H. Toki, and W. Weise, [Phys. Rep.](#) **83**, 281–380 (1982).
- [361] A. Kurkela, K. Rajagopal, and R. Steinhorst, [Phys. Rev. Lett.](#) **132**, 262701 (2024).
- [362] D. H. Rischke and R. D. Pisarski, [arXiv:nucl-th/0004016 \(2000\)](#).
- [363] J. Lynn, I. Tews, J. Carlson, S. Gandolfi, A. Gezerlis, K. Schmidt, and A. Schwenk, [Phys. Rev. Lett.](#) **116**, 062501 (2016).
- [364] D. Lonardonì, I. Tews, S. Gandolfi, and J. Carlson, [Phys. Rev. Res.](#) **2**, 022033 (2020).

- [365] J. Braun, A. Geißel, and B. Schallmo, *SciPost Phys. Core* **7**, 015 (2024).
- [366] F. Di Clemente, A. Drago, and G. Pagliara, *Astrophys. J.* **967**, 159 (2024).
- [367] A. W. Thomas and W. Weise, *The Structure of the Nucleon* (Wiley, Mar. 2001).
- [368] N. Kaiser and W. Weise, [arXiv:2404.11292](https://arxiv.org/abs/2404.11292) (2024).
- [369] Y.-H. Lin, H.-W. Hammer, and U.-G. Meißner, *Phys. Rev. Lett.* **128**, 052002 (2022).
- [370] A. A. Filin, D. Möller, V. Baru, E. Epelbaum, H. Krebs, and P. Reinert, *Phys. Rev. C* **103**, 024313 (2021).
- [371] N. Kaiser and E. Passemar, *Eur. Phys. J. A* **55**, 16 (2019).
- [372] G. Brown, M. Rho, and W. Weise, *Nucl. Phys. A* **454**, 669–690 (1986).
- [373] U.-G. Meissner, N. Kaiser, and W. Weise, *Nucl. Phys. A* **466**, 685–723 (1987).
- [374] R. J. Hill, P. Kammel, W. J. Marciano, and A. Sirlin, *Rep. Prog. Phys.* **81**, 096301 (2018).
- [375] D. E. Kharzeev, *Phys. Rev. D* **104**, 054015 (2021).
- [376] O. Benhar, *Particles* **6**, 611 (2023).
- [377] M. Stephanov, K. Rajagopal, and E. Shuryak, *Phys. Rev. Lett.* **81**, 4816–4819 (1998).
- [378] K. Fukushima, *Phys. Rev. C* **67**, 025203 (2003).
- [379] C. S. Fischer, J. Luecker, and C. A. Welzbacher, *Phys. Rev. D* **90**, 034022 (2014).
- [380] A. Masayuki and Y. Koichi, *Nucl. Phys. A* **504**, 668–684 (1989).
- [381] S. Klimt, M. Lutz, and W. Weise, *Phys. Lett. B* **249**, 386–390 (1990).
- [382] O. Scavenius, A. Mócsy, I. N. Mishustin, and D. H. Rischke, *Phys. Rev. C* **64**, 045202 (2001).
- [383] S. Roessner, C. Ratti, and W. Weise, *Phys. Rev. D* **75**, 034007 (2007).
- [384] S. Roessner, T. Hell, C. Ratti, and W. Weise, *Nucl. Phys. A* **814**, 118–143 (2008).
- [385] T. Hell, S. Rößner, M. Cristoforetti, and W. Weise, *Phys. Rev. D* **81**, 074034 (2010).
- [386] V. Skokov, B. Friman, E. Nakano, K. Redlich, and B.-J. Schaefer, *Phys. Rev. D* **82**, 034029 (2010).
- [387] U. S. Gupta and V. K. Tiwari, *Phys. Rev. D* **85**, 014010 (2012).
- [388] A. Zacchi and J. Schaffner-Bielich, *Phys. Rev. D* **97**, 074011 (2018).
- [389] J. Eser and J.-P. Blaizot, *Phys. Rev. C* **109**, 045201 (2024).
- [390] J. Leong, T. F. Motta, A. W. Thomas, and P. A. M. Guichon, *Phys. Rev. C* **108**, 015804 (2023).
- [391] G. Baym and S. A. Chin, *Nucl. Phys. A* **262**, 527–538 (1976).
- [392] A. Schwenk, G. E. Brown, and B. Friman, *Nucl. Phys. A* **703**, 745–769 (2002).
- [393] G. Baym and C. Pethick, *Landau Fermi-Liquid Theory: Concepts and Applications* (Wiley, Dec. 1991).
- [394] J. Lange, R. O’Shaughnessy, and M. Rizzo, [arXiv:1805.10457](https://arxiv.org/abs/1805.10457) (2018).

- [395] G. Gonçalves, M. Ferreira, J. Aveiro, A. Onofre, F. F. Freitas, C. Providência, and J. A. Font, *JCAP* **12**, 001 (2023).
- [396] F. H. Vivanco, R. Smith, E. Thrane, P. D. Lasky, C. Talbot, and V. Raymond, *Phys. Rev. D* **100**, 103009 (2019).
- [397] J. McGinn, A. Mukherjee, J. Irwin, C. Messenger, M. J. Williams, and I. S. Heng, [arXiv:2403.17462](https://arxiv.org/abs/2403.17462) (2024).
- [398] D. Klochkov, V. Suleimanov, G. Pühlhofer, D. G. Yakovlev, A. Santangelo, and K. Werner, *Astron. Astrophys.* **573**, A53 (2014).
- [399] B. Posselt and G. G. Pavlov, *Astrophys. J.* **932**, 83 (2022).
- [400] L. Lindblom and N. M. Indik, *Phys. Rev. D* **89**, 064003 (2014).
- [401] N. Stergioulas and J. L. Friedman, *Astrophys. J.* **444**, 306 (1995).
- [402] C. Modi, C. Margossian, Y. Yao, R. Gower, D. Blei, and L. Saul, [arXiv:2307.07849](https://arxiv.org/abs/2307.07849) (2023).

