

Data-Driven Generation of Neutron Star Equations of State Using the Variational Autoencoder

Student: Alex Ross¹

Mentors: Tianqi Zhao¹, Sanjay Reddy¹

¹ Institute for Nuclear Theory, University of Washington, Seattle, WA, 98105, USA

January 21, 2026

Abstract

We develop a machine learning model based on the Structured Variational Autoencoder (VAE) framework to reconstruct and generate neutron star (NS) equations of state (EOS). The EOS remains uncertain primarily due to the absence of experimental constraints and the limitations of nuclear theory at supranuclear densities. Our VAE uses an encoder network to map high-dimensional EOS data into a lower-dimensional latent space, capturing essential physical features and NS observables. The corresponding decoder network reconstructs the full EOS from the latent representation with high fidelity. The VAE is trained on a dataset derived from the Skyrme nuclear model. NS observables, including maximum mass and radius obtained by solving the Tolman-Oppenheimer-Volkoff equations, are used as supervised latent parameters. Sampling the latent space allows generation of new, causal and stable EOS models that satisfy known astronomical constraints while exploring regimes beyond existing theories. This framework supports Bayesian inference of the EOS and enables its integration with multimessenger data, including gravitational waves from LIGO/Virgo, X-ray pulsar observations from NICER, and radio pulsar timing measurements of massive NSs. We find that, on the theoretically derived test set, the decoder reconstructs the supervised NS observables with high fidelity, achieving mean absolute percentage errors (MAPE) of 0.10% for the maximum mass M_{\max} and 0.15% for the radius of a $1.4 M_{\odot}$ NS, $R_{1.4}$. Additionally, when analyzing the accuracy of the VAE in mimicking the SLy4 EOS, we find the VAE decodes a value of $M_{\max} = 2.043 \pm 0.002 M_{\odot}$ and $R_{1.4} = 11.715 \pm 0.0176$ km, where the quoted value range correspond to the computed decoder MAPE. This is in agreement to the exact SLy4-computed values of $2.046 M_{\odot}$ and 11.717 km.

1 Introduction

The cold neutron star (NS) equation of state (EOS) describes the relationship between the pressure and the energy density at zero temperature in stable matter within a NS [1]. The EOS is well constrained at low densities near the crust by nuclear theory and experiment, and at asymptotically high densities beyond those realized in NSs by perturbative QCD. However, the intermediate-to-high density EOS relevant to NS cores remains highly uncertain [2] [3] [4]. Matter at extreme densities (up to twice nuclear saturation density) can be described by *ab initio* nuclear theory frameworks such as Chiral Effective Field Theory (χ EFT) [5]. Additionally, phenomenological models such as relativistic mean-field and the Skyrme nuclear models can be used to extrapolate nuclear interactions to higher densities, but these mean field approaches lack higher order quantum corrections to the many-body forces, and the associated theoretical uncertainties are not quantifiable. In this regime, phase transitions may play an important role due to the possible emergence of exotic degrees of freedom (e.g., hyperons or deconfined quarks) [6] [7]. As a result, different theoretical frameworks can produce EOS that are consistent with low-density constraints yet diverge significantly at higher densities, leading to large variations in predicted NS masses, radii, and tidal deformabilities. This uncertainty in the EOS motivates the development of flexible, data-driven approaches that can systematically explore the space of physically admissible EOS while remaining consistent with both theoretical constraints and astrophysical observations. Such data driven approaches have, as of late, largely included machine learning (ML) as a way of rapidly learning, condensing, and scanning a large parameter space.

In particular, using ML for studying and simulating the NS EOS is a subject of recent interest. Many works focus on the use of the Gaussian Process (GP) conditioned on theoretical nuclear models [8] [9] [10]. These studies focus on conditioning the GP on observations to get posterior probability distributions for EOS parameters and observables such as the radius and maximum mass. Works such as [3] and [11] used ML to study the relationship between NS observables and the underlying EOS parameters. In [3], the authors employ the variational autoencoder (VAE) to generate physically valid NS EOS, incorporating observational data to constrain the possible parameter space. The VAE is a deep generative model which is capable of combining the uncertainty of multiple frameworks to constrain the parameter space of the EOS. The primary difference between the GP and the VAE in this context is that the VAE can be used to generate singular EOS by sampling from a probability distribution derived from the encoder, whereas the GP is unable. The work of [3] resulted in the generation of reasonable values for the maximum mass and radius of a $1.4 M_{\odot}$ NS, displaying the validity of the VAE in generating NS EOS. In this work, we expand on this VAE-based approach by systematically analyzing the physical structure and interpretability of the learned latent space, quantifying how individual latent directions map to variations in the decoded EOS and NS observables. We train a fully generative VAE model using data derived from the Skyrme nuclear model, using this input data to calculate theoretical maximum stable NS mass configuration, and the radius of a $1.4 M_{\odot}$ NS which will be used as supervised latent observables during model training. Once we have a fully trained and generative model, we vary parameters to decode the behavior of the NS EOS when controlled by 3 parameters rather than the 10 parameters required by the Skyrme model. By varying these parameters, we can control the behavior of the generated EOS. For example, for a given EOS, scaling the value of the parameter corresponding to the maximum stable mass configuration of a NS should provide a stiffer EOS at high densities. Scaling the parameter corresponding to the radius of a $1.4 M_{\odot}$ NS ($R_{1.4}$) should provide identical EOS for each multiple of $R_{1.4}$ except around twice nuclear saturation density, where $R_{1.4}$ is known to be more sensitive to changes in the EOS [12] and each EOS will differ. We additionally can vary the learned latent

parameter(s) z_i to test and determine which further properties of the EOS is being controlled by said parameter. Once the VAE has been validated, it can be used in future studies as a way to constrain model parameters and simulate EOS rapidly. The paper is organized as follows: in Sec. 2, we discuss the theoretical formulation behind the input data and EOS reconstruction as well as the structure of the VAE framework. In Sec. 3, we test the model as discussed above, analyzing the reconstruction accuracy and mass-radius curve generation capabilities of the framework. In Sec. 4, we summarize our results and in Sec. 5, we offer potential avenues for future research.

2 Methods

2.1 Energy Density Functional

Before we discuss the VAE, it is necessary to formalize the data that the framework will be trained on. The VAE is trained on a diverse candidate EOS dataset derived from the Skyrme nuclear model. The Skyrme model is a phenomenological realization of nuclear density functional theory in which nuclear interactions are encoded in an effective, density-dependent energy density functional rather than derived from the bare nucleon–nucleon force [13]. Originally introduced by Tony Skyrme, the interaction is constructed to reproduce bulk properties of infinite nuclear matter, such as the saturation density, binding energy per nucleon, incompressibility, and symmetry energy, while remaining computationally tractable for finite nuclei and NSs [14]. The resulting Skyrme energy density functional yields self-consistent mean-field equations whose parameters are calibrated to experimental nuclear data, rather than to two-body scattering observables .

The Skyrme Hamiltonian density for infinite nuclear matter is written as follows [15]:

$$H_{\text{Sk}} = \frac{k_{Fn}^5}{10\pi^2 m_n^*} + \frac{k_{Fp}^5}{10\pi^2 m_p^*} + H_{\text{pot}}(n_n, n_p), \quad (1)$$

where the first two terms are the kinetic contributions for neutrons and protons including the effective mass, and H_{pot} is the potential term given by

$$H_{\text{pot}} = \frac{1}{2} n^2 t_0 \left(1 + \frac{x_0}{2}\right) - \frac{1}{2} (n_n^2 + n_p^2) t_0 \left(\frac{1}{2} + x_0\right) + \frac{1}{24} n^\gamma t_3 [n^2(2 + x_3) - (n_n^2 + n_p^2)(1 + 2x_3)]. \quad (2)$$

The neutron and proton effective masses are density dependent and given by

$$\frac{m_{n/p}}{m_{n/p}^*} = 1 + \frac{m_{n/p}}{4} \left\{ n [t_1(2 + x_1) + t_2(2 + x_2)] + n_{n/p} [-t_1(1 + 2x_1) + t_2(1 + 2x_2)] \right\}. \quad (3)$$

Although highly successful in describing ground-state properties and collective excitations of finite nuclei, the Skyrme functional is largely empirical and its extrapolation to the high densities encountered in NS interiors is uncertain, motivating complementary relativistic and data-driven approaches [16] [17].

Table 1: VAE Boundary Value Data

Core–Crust Transition	Maximum
Baryon number density	$n_{B,cc}$
Energy density	ε_{cc}
Pressure	P_{cc}
	n_{max}
	ε_{max}
	P_{max}

2.2 Training Data

We use the Skyrme model discussed in Sec. 2.1 to derive our input EOS training data. The structure of the input data is an array of dimensions $[\mathbf{m}, 107]$, where m is the number of distinct EOS. The first 101 columns are comprised of sound speed data, with the sound speed c_s^2 given by the following equation:

$$c_s^2(\zeta) = \frac{dp}{d\varepsilon}, \quad \zeta \equiv \ln p, \quad (4)$$

The final 6 columns that comprise the input data are boundary conditions that set critical components of the NS structure. This includes theoretical maximum limits of baryon number, energy density, and pressure at the core of a stable NS, as well as the baryon number, energy density, and pressure at the core–crust transition. These core–crust transition values are used to smoothly join the target core EOS to the crust EOS predicted by the Skyrme nuclear model. The boundary value data used is summarized in Tab. 1. A discussion of how they are calculated is included later in Sec. 2.5. The entire input data array is then log-scaled to stabilize the training process. To acquire our pressure values, we define a logarithmically spaced grid as follows;

$$p_i = \exp \left[\ln p_{cc} + \frac{i}{N} (\ln p_{max} - \ln p_{cc}) \right], \quad (5)$$

with $N = 101$ points spanning the interval between the core–crust transition pressure P_{cc} and the pressure of a stable maximum-mass configuration NS P_{max} . The Skyrme nuclear model is used to calculate P_{cc} and P_{max} . For each pressure value P_i , the corresponding energy density $\varepsilon(P_i)$ is obtained from the SLy4 EOS table (a specific parametrization of the Skyrme effective nucleon-nucleon interaction [18] [19]). We then numerically compute the sound speed via a finite-difference scheme using Eq. (4). This data is then split into smaller datasets that are used for different purposes during training. The training split is summarized in Tab. 3. The bulk of the data is allocated to the training dataset, which is used during to minimize the loss function during the VAE training process. The validation dataset serves to tune hyperparameters and decide when training should be terminated without biasing the final result. The test dataset is not used at all during training to calculate gradients or update weightings, and is necessary to evaluate final model performance and report unbiased reconstruction accuracy.

2.3 Reconstructing EOS from Data

Our primary goal in this work is to create new candidate EOS using sound speed profiles predicted by our VAE. We use the boundary values discussed in Sec. 2.2 to define the physically valid range of each EOS sample, allowing the VAE to learn only within the causal and thermodynamically stable region between the crust–core transition and the maximum central density. To calculate the NS EOS with the sound speed, we use the following formulation. Starting from Eq. (4), we note that

$$p = e^\zeta, \quad dp = e^\zeta d\zeta. \quad (6)$$

Substituting (6) into (4) gives

$$d\varepsilon = \frac{dp}{c_s^2(\zeta)} = \frac{e^\zeta}{c_s^2(\zeta)} d\zeta. \quad (7)$$

Integrating from a reference state (p_0, ε_0) to (p, ε) ,

$$\varepsilon(p) = \varepsilon_0 + \int_{\ln p_0}^{\ln p} \frac{e^\zeta}{c_s^2(\zeta)} d\zeta, \quad (8)$$

which is equivalently, in pressure space,

$$\varepsilon(p) = \varepsilon_0 + \int_{p_0}^p \frac{1}{c_s^2(\ln p')} dp'. \quad (9)$$

Where $\epsilon(p)$ is the EOS expressed in terms of the energy density as a function of pressure. In general, Eqns. (8) and (9) have no closed form unless c_s^2 has a simple analytic expression, otherwise they should be evaluated numerically.

The input training data and output VAE-generated data consists of a set of sampled sound-speed values $\{c_s^2(p_i)\}_{i=1}^{101}$, where p_i is calculated using 5, together with thermodynamic boundary conditions at the core-crust transition and at the maximum density. These sound speed values are spline-interpolated to obtain a continuous function $c_s^2(\ln p)$. We then solve Eq. (9) using this function to acquire the corresponding $\epsilon(p)$. The baryon number density $n_B(p)$ is obtained by integrating the ODE

$$\frac{d \ln n_B}{d \ln p} = \frac{1}{\epsilon + p} \frac{d \epsilon}{d \ln p}. \quad (10)$$

The resulting EOS segment $\{n_B(p), \epsilon(p), p\}$ is then joined to the SLy4 crust EOS.

2.4 The TOV Equations

The TOV equations describe the structure of a static NS in hydrodynamical equilibrium [20] [21] [22]. The equations require an input EOS to solve, unless reformulated in a dimensionless form, where the pressure and energy density are expressed as polynomial functions of a reduced radial coordinate as in [20]. However, employing an explicit input EOS remains the standard approach for modeling NS structure. The TOV equations for a spherically symmetric body in static equilibrium are given by,

$$\begin{aligned} \frac{dp(r)}{dr} = & -\frac{G \epsilon(r) m(r)}{r^2} \left(1 + \frac{p(r)}{\epsilon(r) c^2}\right) \left(1 + \frac{4\pi r^3 p(r)}{m(r) c^2}\right) \\ & \times \left(1 - \frac{2Gm(r)}{rc^2}\right)^{-1} \end{aligned} \quad (11)$$

and,

$$\frac{dm(r)}{dr} = 4\pi r^2 \frac{\epsilon(r)}{c^2}, \quad (12)$$

where $\epsilon(r)$ is the energy density as a function of pressure, and $p(r)$ is the pressure as a function of radius. In this work, we solve the TOV equations at two separate stages, prior to and post-training. Prior to training, we use input EOS as calculated using the Skyrme model to compute theoretically possible maximum NS masses (M_{max} , in M_\odot) and radii of 1.4 M_\odot NS ($R_{1.4}$, in km)

with the TOV equations, which are used as supervised latent observables during training. We then solve the equations again post-training, using the sound speed profile $c_s^2(p)$ generated by the VAE to solve Eq. (9) and using the corresponding EOS as input to the TOV equations, to again calculate M_{max} and $R_{1.4}$ and create MR curves for the generated EOS. To create these curves, we pick a central density and calculate the central pressure for a given EOS. Our TOV solver then uses the SciPy LSODA method [23] to integrate Eq. (11) and Eq. (12) outward from this central pressure until the surface pressure approaches zero, yielding the total mass and radius. Repeating this process across a range of central pressures produces the full MR curve predicted by a given generated EOS. Once the high-density portion is computed, it is smoothly joined to the SLy4 crust EOS.

2.5 Boundary Condition Computation

We now discuss how we compute the boundary conditions mentioned in Sec. 2.2 and in Tab. 1. These values are computed using our input EOS as part of our initial data using the Skyrme model, as well as after training using the VAE generated EOS. The core-crust transition baryon number density $n_{B,cc}$ is the density at which uniform nuclear matter becomes unstable to fluctuations in density. Below $n_{B,cc}$, phase separation occurs. The determinant of the Hessian matrix,

$$\det[H] = \frac{\partial^2 E(n_B, x)}{\partial n_B^2} \frac{\partial^2 E(n_B, x)}{\partial x^2} - \left[\frac{\partial^2 E(n_B, x)}{\partial n_B \partial x} \right]^2 > 0 \quad (13)$$

is the stability limit of phase separation, where x is the proton fraction. The total energy per baryon is

$$E(n_B, x) = \varepsilon_N(n_B, x)/n_B + E_e(n_B, x), \quad (14)$$

where ε_N , the energy density of nuclear matter, is given by the Skyrme Hamiltonian density (Eq. (1)),

$$\varepsilon_N = H_{Sk} = \frac{k_{Fn}^5}{10\pi^2 m_n^*} + \frac{k_{Fp}^5}{10\pi^2 m_p^*} + H_{pot}(n_n, n_p), \quad (15)$$

with $n_n = (1 - x)n_B$ and $n_p = xn_B$, and E_e is the energy of a relativistic electron gas with $n_e = xn_B$ (charge neutrality):

$$E_e(n_e) = \frac{1}{\pi^2 n_e} \int_0^{k_{Fe}} \sqrt{k^2 + m_e^2} k^2 dk \approx \frac{3}{5} k_{Fe}, \quad k_{Fe} = (3\pi^2 n_e)^{1/3}. \quad (16)$$

$n_{B,cc}$ is obtained by solving the thermodynamic spinodal condition $\det[H] = 0$ along the beta-equilibrated EOS where the second term in Eq (13) vanishes, $\frac{\partial E(n_B, x)}{\partial x} = 0$. This condition marks the onset of an instability of uniform nuclear matter toward clustering into nuclei.

To calculate n_{max} , or the baryon number density of the maximum stable NS mass configuration, we first solve the TOV equations as given in Sec. 2.4 to acquire the maximum stable mass point. We then evaluate the n_{max} at this point and use it as part of our boundary condition dataset. We also take the pressure P_{max} and energy density ε_{max} at this point to fully constrain the upper boundary on our EOS. To compute the pressure and energy density at the core-crust transition density (P_{cc} and ε_{cc} , respectively), we simply evaluate the EOS at the density $n_{B,cc}$.

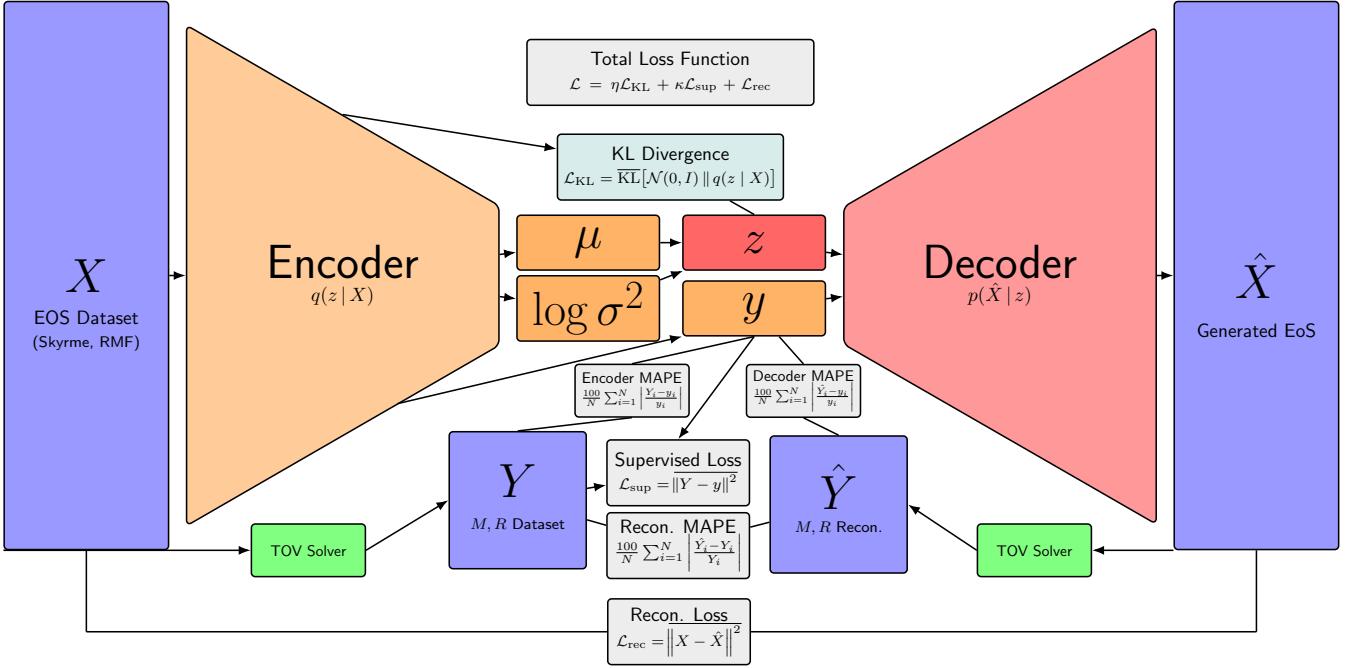


Figure 1: The Variational Autoencoder Framework.

2.6 The Structured Variational Autoencoder

The VAE is an artificial neural network architecture that consists of the components shown in Fig. 1 [24] [25] [26]. The neural network architecture and activation function at each layer used in this work is given in Tab. 2. The goal of the encoder is to learn the probability distribution of the lower-dimension latent variable z given the input dataset X discussed in Sec. 2.2. This dimensionality reduction enables the network to compress complex, high-dimensional EOS information into a smooth latent manifold that captures the underlying physical correlations and variability in data. In doing so, it regularizes the learning process, mitigates overfitting to numerical artifacts, and facilitates an interpretable mapping between microphysical features of the EOS and macroscopic NS observables such as M_{\max} and $R_{1.4}$. A standard autoencoder directly outputs latent vectors from the encoder, while the VAE has a sampling layer composed of the mean μ and the log of the variance $\log(\sigma^2)$ calculated from the encoder. From this sampling layer, we draw a random sample $\epsilon \sim \mathcal{N}(0, I)$ from the standard normal distribution and compute latent vectors using the transformation $z = \mu + \exp(0.5 \cdot \log \sigma^2) \cdot \epsilon$. This formulation enables stochastic sampling from the latent Gaussian distribution while preserving differentiability with respect to μ and σ . Consequently, gradients can propagate through the sampling process during backpropagation, allowing the encoder to learn an appropriate latent distribution.

Additionally, the encoder learns a distribution y of the supervised latent observables (M_{\max} and $R_{1.4}$) as a part of the latent layer for later use. The sampled latent vectors z are then passed into the decoder, which increases dimensionality and computes a new probability distribution \hat{X} of z . The output of the decoder is the reconstructed c_s^2 data with boundary conditions, which is used to compute the generated EOS, and is then passed into an algorithm that computes mass-radius (MR) curves for a given EOS using the TOV equations from Sec. 2.4. These MR values are stored in \hat{Y} .

Throughout training, a total loss function is constructed to jointly optimize the reconstruction accuracy of the EOS, the supervised agreement between the predicted and true supervised latent

observable values, and the Kullback–Leibler divergence that regularizes the latent space toward a standard normal prior, thereby balancing physical fidelity, predictive performance, and latent-space smoothness. The loss function used is given by

$$\mathcal{L} = \mathcal{L}_{rec} + \kappa \mathcal{L}_{sup} + \eta \mathcal{L}_{KL} \quad (17)$$

$$\begin{aligned} \mathcal{L} = & \overline{\|X - \hat{X}\|^2} + \kappa \overline{\|Y - y\|^2} \\ & + \frac{\eta}{N \cdot d} \sum_{n=1}^N \sum_{i=1}^d \left[-\frac{1}{2} (1 + \log \sigma_{n,i}^2 - \mu_{n,i}^2 - \sigma_{n,i}^2) \right] \end{aligned} \quad (18)$$

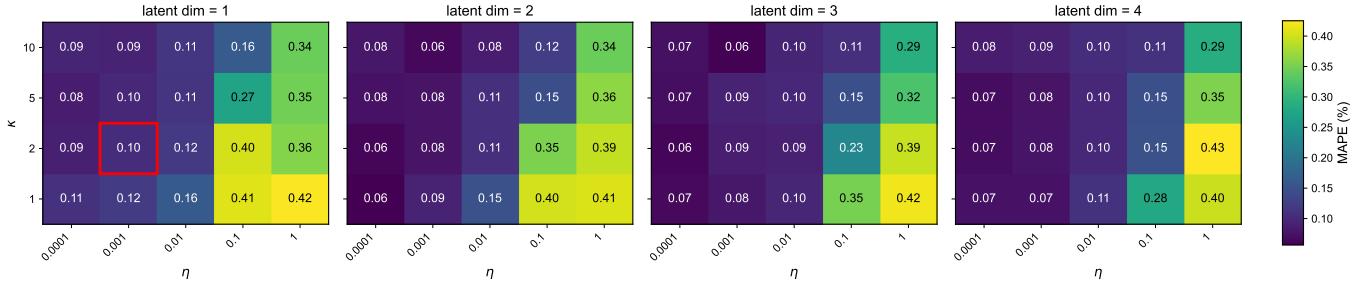
Where \mathcal{L}_{rec} is the reconstruction loss between the input and reconstructed data, \mathcal{L}_{sup} is the reconstruction loss between the input supervised latent observables and those predicted by the encoder, \mathcal{L}_{KL} is the Kullback–Leibler (KL) divergence, and η, κ are weighting terms. The KL divergence is a measure of how different our prior $\sim \mathcal{N}(0, I)$ is from our probability distribution $q(z|X)$ predicted by the encoder. We compute it for each latent dimension i up to the total number of latent dimensions d , sum the contribution of each dimension, and take the mean by dividing by the number of samples in the batch N multiplied by the number of latent dimensions. The KL divergence acts as a regularizing term, meaning that it can help coax the probability distribution of the variational latent variables towards a normal gaussian profile by penalizing it for deviations [3]. Thus, a higher KL weighting η forces the latent space distribution closer to a normal distribution. η must be selected to be sufficiently small in order to minimize MAPE while avoiding posterior collapse, in which the latent variables become weakly informative and the decoder ignores the latent structure. Additionally, as η goes to zero, we recover an autoencoder framework as we lose the probabilistic sampling nature of the VAE [27].

Table 2: Neural Network Architecture

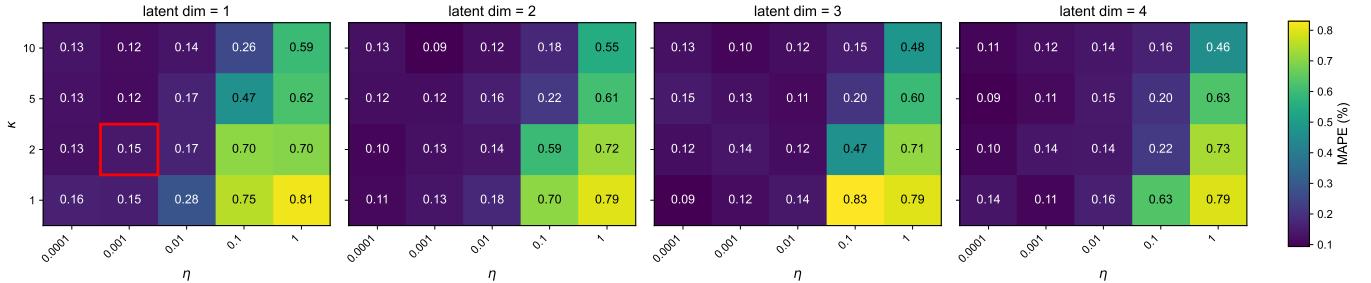
Layer	Type	Neurons	Activation
Input	—	107	—
Layer 1	Dense	64	ReLU
Layer 2	Dense	64	ReLU
Latent layer	Lambda	4	—
Layer 3	Dense	64	ReLU
Layer 4	Dense	64	ReLU
Output	Dense	107	—

As outlined in Tab. 2, the ReLU activation function is used in our work. The ReLU (Rectified Linear Unit) activation function allows positive values to pass through unmodified, while clipping negative values to zero [28]. The output activation function choice can be better selected to ensure causality, perhaps the sigmoid function. An activation function for the output layer did not affect our results of this work, but could be an interesting area of further testing and exploration.

The primary machine learning system we use is the Python package Tensorflow [29]. We use a batch size of 64 and a learning rate of $\alpha = .0001$. Our input data consisting of the array discussed in Sec. 2.2, as well as our supervised latent observables discussed in 2.4, undergo a transformation using the StandardScaler given by Scikit-learn. This standardizes the data by removing the mean and scaling to unit variance. For the selected dimension of the latent space, we refer to the mean absolute percentage error (MAPE) heatmap shown in Fig. 2. The MAPE is computed by the equations



(a) MAPE heatmap for the maximum mass M_{\max}



(b) MAPE heatmap for the radius of a $1.4 M_{\odot}$ NS ($R_{1.4}$)

Figure 2: Decoder MAPE comparison for both supervised latent observables defined in Fig. 1. The MAPE is computed for each combination of latent dimensionality, κ , and η used during training. The hyperparameter and dimensionality combo selected for further analysis is outlined with a red box.

shown in Fig. 1. The heatmap shows the MAPE computed between the decoder-predicted values of our supervised latent observables and the encoder-learned supervised latent observables that are part of the test set. This MAPE is particularly important since it measures how faithfully the decoder preserves the mapping between the latent variables and physically observable NS properties. A low decoder MAPE therefore indicates that the latent space encodes sufficient information to accurately reproduce the target observables, although it does not by itself imply physical interpretability or disentanglement of individual latent dimensions. We train a model and compute the MAPE for each shown combination of hyperparameter η and κ , and for latent dimensions 1-4. Each tile in the heatmap represents a different training run using the corresponding hyperparameters and dimensionality. We use this heatmap to select the model that we will analyze further in Sec. 3. We choose a combination with a low MAPE that is similar for both supervised latent observables, ensuring that the latent space encodes both quantities with similar fidelity. For the MAPE corresponding to the M_{\max} , a value of .10%, and .15% for the value corresponding to $R_{1.4}$. This corresponds to hyperparameter values of $\eta = 0.001$ and $\kappa = 2$, with a latent dimensionality of one. The selected combination is outlined with a red box in Fig. 2. The latent dimension choice can be made by inspecting the latent space, as we do in Sec. 3.

3 Results

After training is completed, we now have a fully generative VAE model that we can sample from to generate EOS. The EOS can be fully controlled using a single latent parameter z_0 . The decoder,

Table 3: Dataset Splits Used for Training and Evaluation

Split	Size	Purpose	During Training
Training	~70%	Fit model parameters	Yes — updates weights
Validation	~15%	Hyperparameter and epoch selection	No — evaluation only
Test	~15%	Final unbiased performance	No — never used during training

conditioned by the supervised latent observables M_{max} and $R_{1.4}$, is able to make sense of the single latent parameter and convert z_0 into meaningful EOS. Thus, we can now mimic a ten-parameter Skyrme model using only two supervised latent observables and one latent variable, for a total of three parameters. In this section, we will analyze the model indicated in Fig. 2 in detail.

3.1 Latent Space

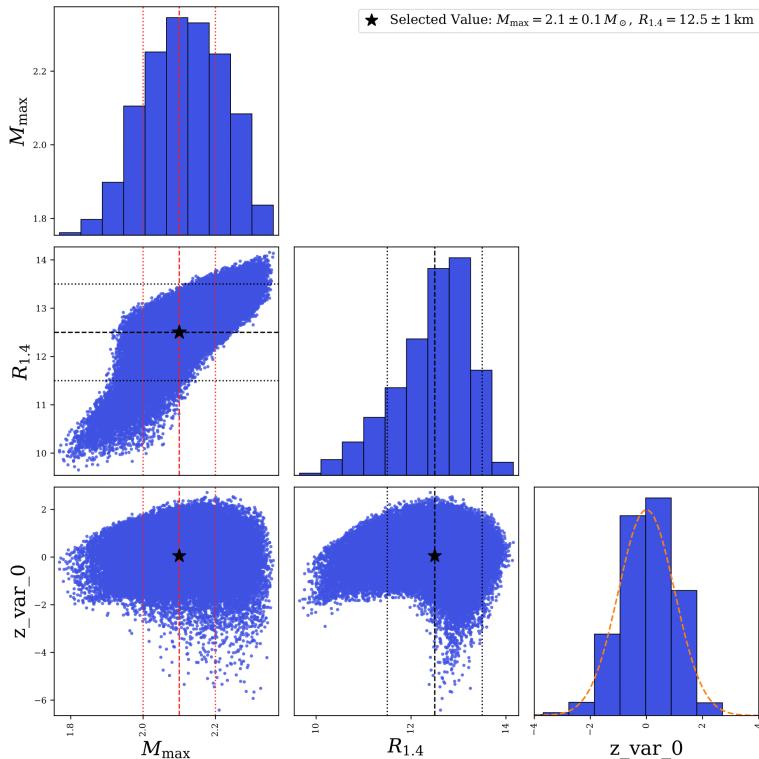


Figure 3: Pairwise distributions for the supervised and latent variables for test data. The selected central values of M_{max} and $R_{1.4}$, along with the intervals used to probe latent space sensitivity in the reconstructed EOS, are indicated. For the latent variable z_0 , a standard normal distribution (orange) is overlaid on the histogram.

Fig. 3 shows the latent space distribution of the test dataset, corresponding to the parameter space that the decoder can sample from. This trained VAE model has one latent dimension and two supervised latent observables, M_{max} and $R_{1.4}$, with hyperparameters $\eta = 0.001$ and $\kappa = 2$. This configuration was selected based on the criteria discussed in Sec. 2.6. We choose the number of latent variables z_i based on the minimum dimensionality required to capture nontrivial, physically meaningful variability in the EOS. Latent dimensions that collapse to near-linear manifolds or exhibit

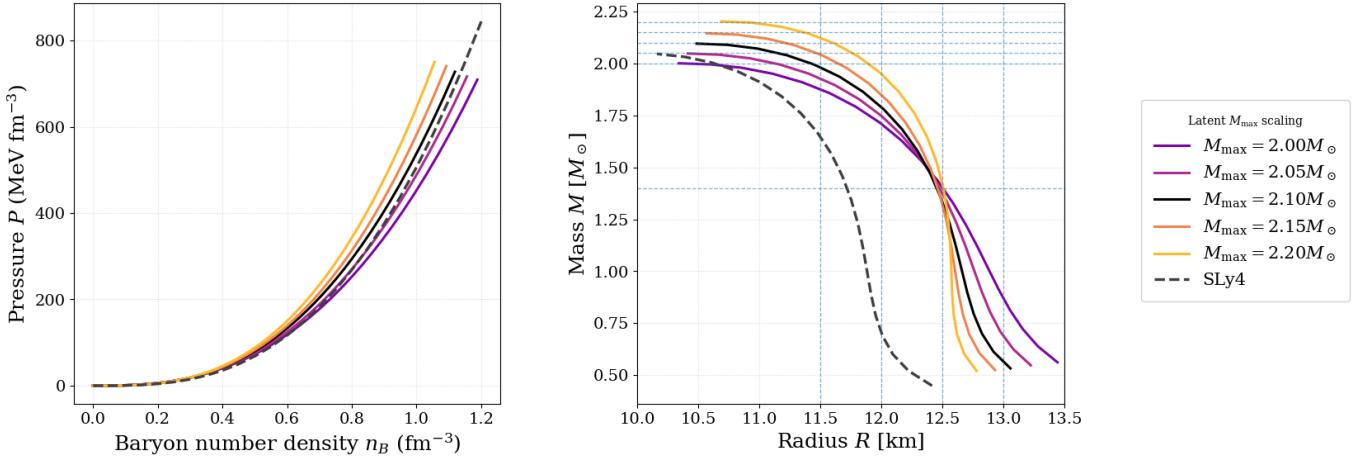


Figure 4: EOS $P(n_B)$ generated by varying the supervised latent parameter M_{\max} about a central value of $2.1 M_{\odot}$, while holding $R_{1.4}$ and z_0 fixed. For each chosen M_{\max} , the corresponding EOS is decoded from the VAE latent space and evaluated over a common pressure grid. For comparison, we also show the SLy4 EOS (dashed). The right panel shows the mass–radius curves computed from each corresponding EOS.

strongly non-Gaussian, degenerate distributions are indicative of redundant or inactive modes and are therefore excluded. The supervised latent observables exhibit a strong linear correlation as expected, quantified by the Pearson coefficient $r(M_{\max}, R_{1.4}) = 0.8117$, computed over the full test set. The latent variable z_0 in Fig. 3 follows an approximately Gaussian distribution as imposed by the KL divergence. When projected against the supervised latent observables, we see a smooth, nonlinear correlation, indicating that this latent dimension has learned a physically meaningful mode of variation in the EOS.

3.2 Decoded EOS and Mass-Radius Curves

To test the generative capabilities of the trained VAE model, we vary individual latent parameters and supervised latent observables over their posterior-supported ranges while fixing all other dimensions at their mean values. Each perturbation is decoded into a full EOS, which is then propagated through the TOV equations to produce MR relations. This procedure allows us to verify that smooth, physically interpretable variations in latent space correspond to continuous, stable, and astrophysically consistent deformations of NS structure, thereby validating both the physical expressivity and generative robustness of the learned model.

Fig. 4 shows the first such perturbation. We fix the latent variable z_0 and the supervised latent observable $R_{1.4}$ while varying the other supervised latent observable, M_{\max} about the central selected value as indicated in Fig. 3. For the EOS $P(n_B)$, the decoded equations of state remain tightly clustered at low baryon densities, indicating that the crust and outer-core behavior is largely unaffected by changes in the supported maximum mass. At higher densities, however, the EOS curves begin to separate noticeably, reflecting increasing sensitivity to the assumed high-density physics. We find that the decoded equations of state become systematically stiffer at high densities as M_{\max} is increased, consistent with the physical expectation that greater pressure support at supranuclear densities is required to stabilize more massive NSs against gravitational collapse [30]. The exact SLy4 EOS and corresponding MR curve are plotted for reference.

In Fig. 5, we perform the same procedure as in Fig. 4, but we instead fix M_{\max} and z_0 while

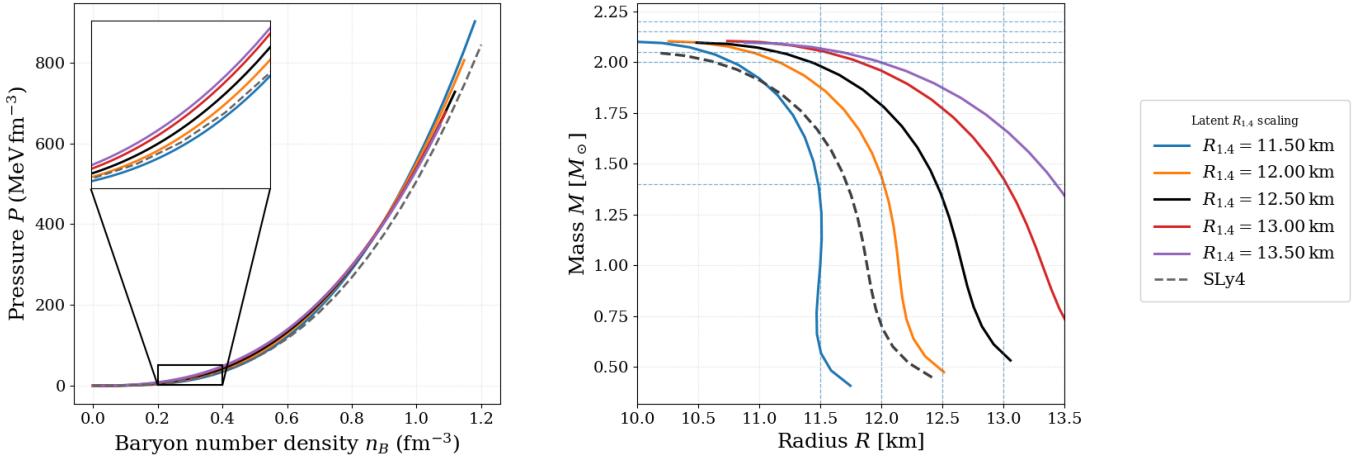


Figure 5: Same as Fig. 4, except the supervised latent parameter M_{\max} and the latent variable z_0 are held fixed while varying $R_{1.4}$ about a central value of 12.5 km.

varying $R_{1.4}$ (the supervised latent observable corresponding to the radius of a $1.4 M_{\odot}$ NS). The EOS remain coherent at higher densities corresponding to the core of a NS, but noticeably diverge at the low-density region highlighted in the upper-left panel. This behavior is consistent with the well-known sensitivity of NS radii to the pressure near and just above nuclear saturation density, where variations in the EOS primarily influence the stellar envelope rather than the core [31].

We do the same procedure as above in Fig. 6, but instead varying the latent variable z_0 while fixing M_{\max} and $R_{1.4}$ and decoding the corresponding EOS. This variation produces EOS that are very similar, with variations at baryon density $n_B \lesssim 0.1 \text{ fm}^{-3}$. Differences in the low-density EOS directly translate into variations in neutron-star radii, particularly for lower-mass stars, indicating that z_0 primarily encodes information relevant to the intermediate-to-low density structure of NS.

An interesting feature to note is the EOS crossing point that occurs at low densities around nuclear saturation density in Fig. 4 (more apparent in log-scaled space), and the corresponding MR curve crossing that occurs around a mass of $1.4 M_{\odot}$ and a radius of 12.5 km. These features are likely correlated and reflect the sensitivity of $R_{1.4}$ to the EOS at near saturation density discussed above. This crossing is clear to see in the EOS curves as shown in Fig. 5.

3.3 SLy4 EOS Comparison

We now aim to show that the VAE framework is capable of accurately reproducing the SLy4 EOS by directly selecting and passing to the decoder the supervised latent observables $M_{\max} = 2.046 M_{\odot}$ and $R_{1.4} = 11.717 \text{ km}$, values calculated with the SLy4 parameterization. Fig. 7 shows the EOS and corresponding MR curves for SLy4, and the VAE-generated EOS made by explicitly passing in the above values of M_{\odot} and $R_{1.4}$ as conditioning points to the decoder. For the value of the latent parameter z_0 , we use the scaled latent mean value of $z_0 = .0487$. We also use the computed value of the decoder MAPE for $R_{1.4}$ of 0.15% (as given in Fig. 2) to construct a symmetric fractional tolerance band about the VAE-generated MR curve. This acts as a visual indicator of the high accuracy the VAE achieves when attempting to decode the SLy4 EOS. In particular, the VAE decodes a value of $M_{\max} = 2.043 \pm 0.002 M_{\odot}$ and $R_{1.4} = 11.715 \pm .0176 \text{ km}$, in very close agreement to the exact SLy4-calculated values of $2.046 M_{\odot}$ and 11.717 km . We scale up the tolerance band in 7 by a factor of 10 for visibility.

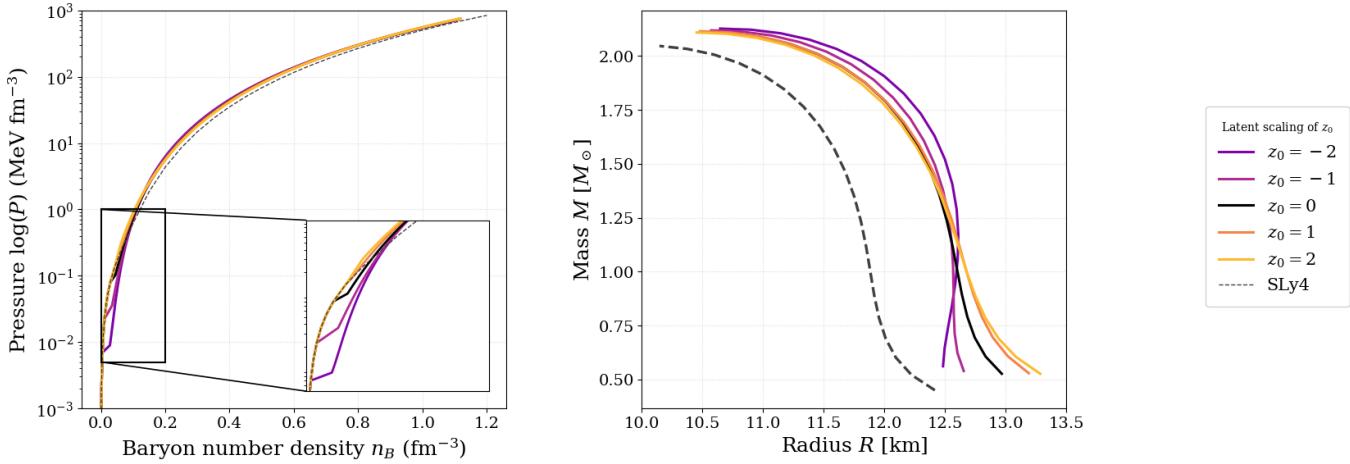


Figure 6: EOS $P(n_B)$ generated by fixing the supervised latent parameters M_{\max} and $R_{1.4}$ and varying the latent parameter z_0 . The left plot is log scaled to highlight the low-density EOS variation, and the SLy4 curve’s appearance has been modified to make the variation visible.

4 Conclusion

In this work, we create a VAE framework to generate new candidate NS EOS. We show that our model is capable of taking theoretical 10 parameter Skyrme EOS as input EOS data, compressing it down to a single latent dimension z_0 and two supervised latent observables M_{\max} and $R_{1.4}$. We can then use this reduced dimension latent space to control the generation of physically consistent NS EOS. By varying the latent space values, we demonstrated that our decoded EOS remain smooth, causal, and thermodynamically stable across the full density range of interest, while exhibiting controlled and physically interpretable variations in both macroscopic NS observables and high-density EOS behavior. In particular, systematic changes in the supervised and variational latent parameters produce coherent shifts in the stiffness of the EOS at different density ranges and the corresponding mass-radius relations at different mass ranges. Together, these results demonstrate that the latent representation learned by the VAE provides a compact and physically meaningful parameterization of the NS EOS, enabling efficient exploration of EOS sensitivity within observationally and theoretically admissible bounds.

5 Future Work

Future work should include a comparison of the speed increases of this VAE approach as compared to traditional methods of creating new EOS, to cement it’s robustness as an alternative, if not improved, method for NS EOS studies. One of the primary advantages of using a decoder network to perform Markov chain Monte Carlo sampling of latent space parameters for Bayesian analysis involving NS observations is the decoder can generate large numbers of candidate EOS realizations essentially “on demand”. This potentially enables faster exploration of EOS uncertainty while maintaining physically motivated constraints embedded in the training data and EOS construction procedure, which needs to be explored in future works.

This framework supports the integration of other theoretical nuclear models into the input data used to train the model. Naturally, adding EOS and boundary conditions derived from models other than Skyrme, such as the Relativistic Mean-Field model, would potentially allow a more thorough

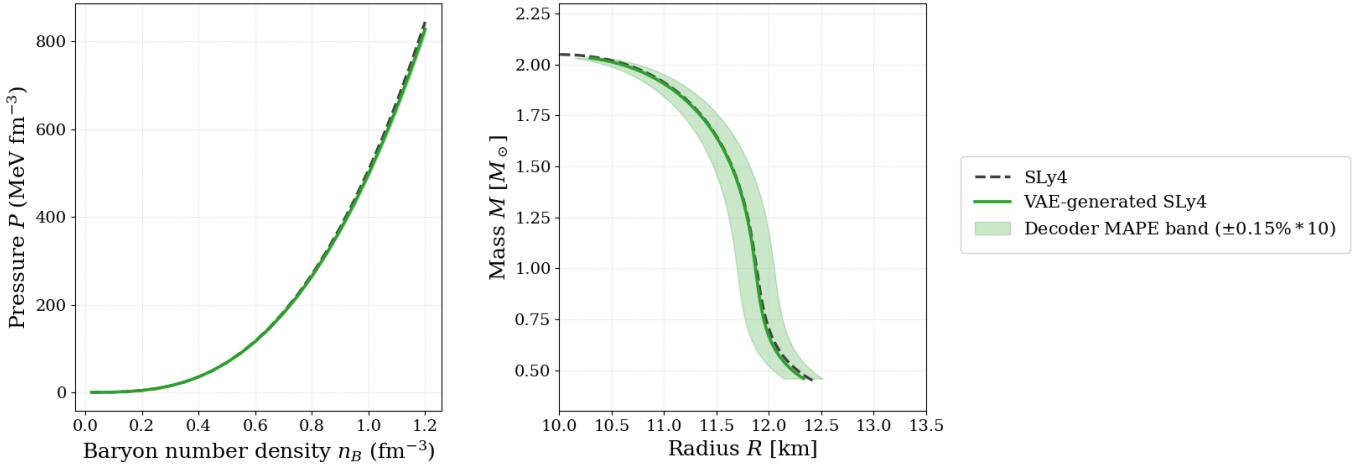


Figure 7: Equation of state $P(n_B)$ and MR curves generated by fixing the values of M_{\max} and $R_{1.4}$ to the SLy4 values of $2.046 M_{\odot}$ and 11.717 km, respectively and decoding the EOS. Shaded regions denote the fractional tolerance band obtained by scaling the decoded EOS by the decoder MAPE of $R_{1.4}$. The exact SLy4 model EOS and MR curve is shown for reference.

exploration of the theoretical EOS parameter space. In addition, scenarios involving first-order phase transitions, in contrast to the relatively smooth EOS functions considered here, should be explored. Such transitions under Maxwell construction between various quark models such as MIT bag model and hadronic model EOS, naturally allow the sound speed to approach zero and introduce discontinuities in the EOS [32] [33]. These more sophisticated possibilities may be particularly well-suited for application of the VAE framework proposed in this work.

Acknowledgement

This research is supported by N3AS’s National Science Foundation award No. 2020275, and a scholarship from the Mary Gates Endowment for Students. We thank the Institute for Nuclear Theory at the University of Washington for its support.

References

- [1] Kohsuke Sumiyoshi, Toru Kojo, and Shun Furusawa. “Equation of state in neutron stars and supernovae”. In: *arXiv e-prints*, arXiv:2207.00033 (June 2022), arXiv:2207.00033. DOI: 10.48550/arXiv.2207.00033. arXiv: 2207.00033 [nucl-th].
- [2] C. Drischler, K. Hebeler, and A. Schwenk. “Chiral interactions up to next-to-next-to-next-to-leading order and nuclear saturation”. In: *arXiv e-prints*, arXiv:1710.08220 (Oct. 2017), arXiv:1710.08220. DOI: 10.48550/arXiv.1710.08220. arXiv: 1710.08220 [nucl-th].
- [3] Ming-Zhe Han, Shao-Peng Tang, and Yi-Zhong Fan. “Nonparametric Representation of Neutron Star Equation of State Using Variational Autoencoder”. In: *The Astrophysical Journal* 950.2, 77 (June 2023), p. 77. DOI: 10.3847/1538-4357/acd050. arXiv: 2205.03855 [astro-ph.HE].
- [4] Gordon Baym et al. “From hadrons to quarks in neutron stars: a review”. In: *RPP* 81.5, 056902 (May 2018), p. 056902. DOI: 10.1088/1361-6633/aaae14. arXiv: 1707.04966 [astro-ph.HE].

[5] C. Drischler, J. W. Holt, and C. Wellenhofer. “Chiral Effective Field Theory and the High-Density Nuclear Equation of State”. In: *Annual Review of Nuclear and Particle Science* 71 (Sept. 2021), pp. 403–432. DOI: 10.1146/annurev-nucl-102419-041903. arXiv: 2101.01709 [nucl-th].

[6] M. Dutra et al. “Relativistic mean-field hadronic models under nuclear matter constraints”. In: *Physics Review C* 90.5, 055203 (Nov. 2014), p. 055203. DOI: 10.1103/PhysRevC.90.055203. arXiv: 1405.3633 [nucl-th].

[7] Chun Huang et al. “Constraining a relativistic mean field model using neutron star mass-radius measurements I: nucleonic models”. In: *Monthly Notices of the Royal Astronomical Society* 529.4 (Apr. 2024), pp. 4650–4665. DOI: 10.1093/mnras/stae844. arXiv: 2303.17518 [astro-ph.HE].

[8] Philippe Landry and Reed Essick. “Nonparametric inference of the neutron star equation of state from gravitational wave observations”. In: *Physics Review D* 99.8, 084049 (Apr. 2019), p. 084049. DOI: 10.1103/PhysRevD.99.084049. arXiv: 1811.12529 [gr-qc].

[9] Wenjie Zhou et al. “Nonparametric Model for the Equations of State of a Neutron Star from Deep Neural Network”. In: *The Astrophysical Journal* 950.2, 186 (June 2023), p. 186. DOI: 10.3847/1538-4357/acd335. arXiv: 2305.03323 [nucl-th].

[10] Debora Mroczek et al. “Searching for phase transitions in neutron stars with modified Gaussian processes”. In: *arXiv e-prints*, arXiv:2302.07978 (Feb. 2023), arXiv:2302.07978. DOI: 10.48550/arXiv.2302.07978. arXiv: 2302.07978 [astro-ph.HE].

[11] N. K. Patra et al. “Inferring the equation of state from neutron star observables via machine learning”. In: *Physics Letters B* 865, 139470 (June 2025), p. 139470. DOI: 10.1016/j.physletb.2025.139470. arXiv: 2502.20226 [nucl-th].

[12] Jeremy W. Holt and Yeunhwan Lim. “Dense matter equation of state and neutron star properties from nuclear theory and experiment”. In: *American Institute of Physics Conference Series*. AIP, July 2019, 020019, p. 020019. DOI: 10.1063/1.5117809. arXiv: 1904.11449 [nucl-th].

[13] Junjie Yang and J. Piekarewicz. “Covariant Density Functional Theory in Nuclear Physics and Astrophysics”. In: *Annual Review of Nuclear and Particle Science* 70 (Oct. 2020), pp. 21–41. DOI: 10.1146/annurev-nucl-101918-023608. arXiv: 1912.11112 [nucl-th].

[14] Carlos Naya. “Neutron stars within the Skyrme model”. In: *IJMP E* 28.8, 1930006 (Jan. 2019), p. 1930006. DOI: 10.1142/S0218301319300066. arXiv: 1910.01145 [astro-ph.HE].

[15] Tianqi Zhao et al. “Characterizing the nuclear models informed by PREX and CREX: a view from Bayesian inference”. In: *arXiv e-prints*, arXiv:2406.05267 (June 2024), arXiv:2406.05267. DOI: 10.48550/arXiv.2406.05267. arXiv: 2406.05267 [nucl-th].

[16] Mingya Duan and Michael Urban. “New Skyrme parametrizations to describe finite nuclei and neutron star matter with realistic effective masses”. In: *Physical Review C* 110.6 (Dec. 2024). ISSN: 2469-9993. DOI: 10.1103/physrevc.110.065806. URL: <http://dx.doi.org/10.1103/PhysRevC.110.065806>.

[17] J.R. Stone and P.-G. Reinhard. “The Skyrme interaction in finite nuclei and nuclear matter”. In: *Progress in Particle and Nuclear Physics* 58.2 (Apr. 2007), pp. 587–657. ISSN: 0146-6410. DOI: 10.1016/j.ppnp.2006.07.001. URL: <http://dx.doi.org/10.1016/j.ppnp.2006.07.001>.

[18] Yingxun Zhang et al. “Constraints on nucleon effective mass splitting with heavy ion collisions”. In: *Physics Letters B* 732 (May 2014), pp. 186–190. DOI: 10.1016/j.physletb.2014.03.030. arXiv: 1402.3790 [nucl-th].

[19] Doan Thi Loan et al. “Equation of state of neutron star matter, and the nuclear symmetry energy”. In: *Physics Review C* 83.6, 065809 (June 2011), p. 065809. DOI: 10.1103/PhysRevC.83.065809. arXiv: 1105.5222 [nucl-th].

[20] Bao-Jun Cai and Bao-An Li. “Novel scalings of neutron star properties from analyzing dimensionless Tolman-Oppenheimer-Volkoff equations”. In: *European Physical Journal A* 61.3, 55 (Mar. 2025), p. 55. DOI: 10.1140/epja/s10050-025-01507-7. arXiv: 2501.18676 [astro-ph.HE].

[21] Richard C. Tolman. “Static Solutions of Einstein’s Field Equations for Spheres of Fluid”. In: *Phys. Rev.* 55 (4 Feb. 1939), pp. 364–373. DOI: 10.1103/PhysRev.55.364. URL: <https://link.aps.org/doi/10.1103/PhysRev.55.364>.

[22] J. R. Oppenheimer and G. M. Volkoff. “On Massive Neutron Cores”. In: *Phys. Rev.* 55 (4 Feb. 1939), pp. 374–381. DOI: 10.1103/PhysRev.55.374. URL: <https://link.aps.org/doi/10.1103/PhysRev.55.374>.

[23] Linda R. Petzold. “Automatic Selection of Methods for Solving Stiff and Nonstiff Systems of Ordinary Differential Equations”. In: *SIAM Journal on Scientific and Statistical Computing* 4.1 (1983), pp. 136–148. DOI: 10.1137/0904010.

[24] Diederik P. Kingma and Max Welling. “An Introduction to Variational Autoencoders”. In: *arXiv e-prints*, arXiv:1906.02691 (June 2019), arXiv:1906.02691. DOI: 10.48550/arXiv.1906.02691. arXiv: 1906.02691 [cs.LG].

[25] Carl Doersch. “Tutorial on Variational Autoencoders”. In: *arXiv e-prints*, arXiv:1606.05908 (June 2016), arXiv:1606.05908. DOI: 10.48550/arXiv.1606.05908. arXiv: 1606.05908 [stat.ML].

[26] Márcio Ferreira and Michał Bejger. “Conditional variational autoencoder inference of neutron star equation of state from astrophysical observations”. In: *Physics Review D* 111.2, 023035 (Jan. 2025), p. 023035. DOI: 10.1103/PhysRevD.111.023035. arXiv: 2403.14266 [nucl-th].

[27] Diederik P. Kingma and Max Welling. “An Introduction to Variational Autoencoders”. In: *arXiv e-prints*, arXiv:1906.02691 (June 2019), arXiv:1906.02691. DOI: 10.48550/arXiv.1906.02691. arXiv: 1906.02691 [cs.LG].

[28] Abien Fred Agarap. *Deep Learning using Rectified Linear Units (ReLU)*. 2019. arXiv: 1803.08375 [cs.NE]. URL: <https://arxiv.org/abs/1803.08375>.

[29] Martín Abadi et al. “TensorFlow: A system for large-scale machine learning”. In: *arXiv e-prints*, arXiv:1605.08695 (May 2016), arXiv:1605.08695. DOI: 10.48550/arXiv.1605.08695. arXiv: 1605.08695 [cs.DC].

[30] J. M. Lattimer and M. Prakash. “Neutron Star Structure and the Equation of State”. In: *The Astrophysical Journal* 550.1 (Mar. 2001), pp. 426–442. DOI: 10.1086/319702. arXiv: astro-ph/0002232 [astro-ph].

[31] K. Hebeler et al. “Equation of State and Neutron Star Properties Constrained by Nuclear Physics and Observation”. In: *The Astrophysical Journal* 773.1, 11 (Aug. 2013), p. 11. DOI: 10.1088/0004-637X/773/1/11. arXiv: 1303.4662 [astro-ph.SR].

[32] Luiz L Lopes, Carline Biesdorf, and Débora P Menezes. “Modified MIT bag Models—part I: Thermodynamic consistency, stability windows and symmetry group”. In: *Physica Scripta* 96.6 (Mar. 2021), p. 065303. ISSN: 1402-4896. DOI: 10.1088/1402-4896/abef34. URL: <http://dx.doi.org/10.1088/1402-4896/abef34>.

[33] Sebastian Blacker, Andreas Bauswein, and Stefan Typel. “Exploring thermal effects of the hadron-quark matter transition in neutron star mergers”. In: Physical Review D 108.6 (Sept. 2023). ISSN: 2470-0029. DOI: 10.1103/physrevd.108.063032. URL: <http://dx.doi.org/10.1103/PhysRevD.108.063032>.