How effectively can Neural Posterior Estimation infer the Neutron Star Equation of State?

Valéria Carvalho[1], Márcio Ferreira[1], Constança Providência[1], Michał Bejger[2,3]

[1]CFisUC, Department of Physics, University of Coimbra,[2]Nicolaus Copernicus Astronomical Center, Polish Academy of Sciences,[3]INFN Sezione di Ferrara

Goal: To infer the neutron star equation of state (EoS) from astrophysical observations while rigorously quantifying uncertainties. **Challenge:** Despite significant observational progress, available measurements of neutron star properties—such as mass, radius, and tidal deformability—remain limited in both quantity and precision. This hinders reliable EoS inference using traditional methods. **Our Approach:** We adopt simulation-based inference, specifically Neural Posterior Estimation (NPE), to overcome these limitations and extract robust EoS constraints from realistic, noisy datasets.

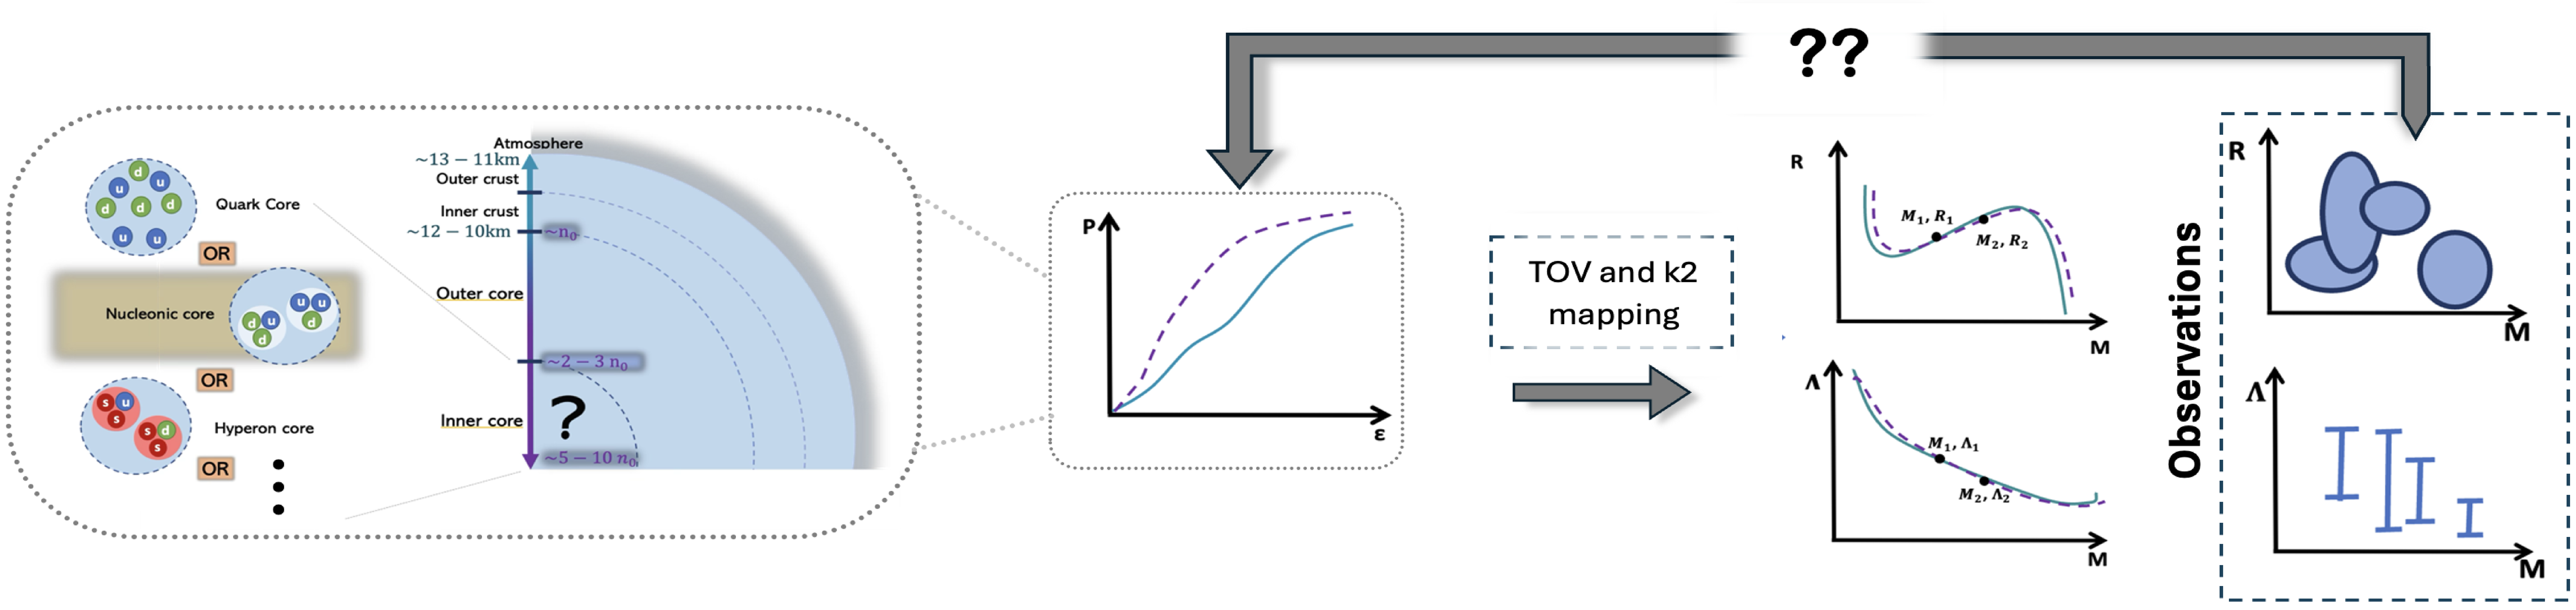

Dataset Generation

We construct synthetic datasets to train a model capable of inferring the neutron star equation of state (EoS) from observational data. Our approach is inspired by the agnostic EoS prior introduced in [1], which allows us to generate physically consistent mock data without committing to a specific microscopic model.

We adapt this prior to define our own simulation-based dataset tailored for training and testing Neural Posterior Estimation (NPE). This setup enables flexibility: from a single base dataset, we derive multiple variants with different levels of observational realism (e.g., with/without noise or tidal deformability).

Pressure Sampling:

$$\boldsymbol{p} = [p(n_1), \dots, p(n_{20})], \quad \text{with } n_1 = 0.13\,\mathrm{fm}^{-3},\ n_{20} = 1.28\,\mathrm{fm}^{-3}.$$

Mass Sampling Strategy: We sample neutron star masses from three astrophysically motivated ranges:

$$M^a \in \mathcal{U}[1.0, 1.4]\,M_\odot, \quad M^b \in \mathcal{U}[1.4, 1.7]\,M_\odot, \quad M^c \in \mathcal{U}[1.7, M_{\text{max}}(\text{EoS})]\,M_\odot$$

Each region contains $n_o = 5$ samples to ensure coverage across low, canonical, and high-mass stars.

Observables: We compute radius and tidal deformability for each sample:

$$\mathbf{M}^0 = [M_1^a, \dots, M_{n_o}^c], \quad \mathbf{R}^0 = [R(M_1^a), \dots, R(M_{n_o}^c)]$$
$$\mathbf{M}^* = [M_1^{*a}, \dots, M_{n_o}^{*c}], \quad \mathbf{\Lambda}^0 = [\Lambda(M_1^{*a}), \dots, \Lambda(M_{n_o}^{*c})]$$

Adding Observational Noise: We perturb true values using Gaussian noise:

$$\mathbf{M} \sim \mathcal{N}(\mathbf{M}^0, \boldsymbol{\sigma_M}^2), \quad \mathbf{R} \sim \mathcal{N}(\mathbf{R}^0, \boldsymbol{\sigma_R}^2), \quad \mathbf{\Lambda} \sim \mathcal{N}(\mathbf{\Lambda}^0, \boldsymbol{\sigma_\Lambda}^2(M^*))$$

Uncertainty Distributions:

$$\boldsymbol{\sigma_M} \sim \mathcal{U}[0, \sigma_M], \quad \boldsymbol{\sigma_R} \sim \mathcal{U}[0, \sigma_R], \quad \boldsymbol{\sigma_\Lambda}(M^*) \sim \mathcal{U}[0, \hat{\sigma}(M^*)]$$

We use:

$$\sigma_M = 0.1\,M_\odot, \quad \sigma_R = 0.3\,\text{km}, \quad \hat{\sigma}(M) = 121483.4 \cdot e^{-5M + 0.37M^2}$$

This setup enables both noise-free and observationally realistic datasets.

Realistic Test Set from Observed Pulsar Masses:

To better reflect astrophysical reality, we constructed a test dataset based on the observed neutron star mass distribution from [2], which compiles 122 pulsar mass measurements (excluding gravitational wave events). The dataset highlights clustering in three characteristic mass intervals: low-mass, canonical, and high-mass.

We maintained the same mass intervals used during training:

$$[1.0, 1.4], \quad [1.4, 1.7], \quad [1.7, M_{\text{max}}(\text{EoS})]\,M_\odot,$$

but adjusted the number of samples to reflect the empirical distribution, assigning 6, 4, and 5 samples per interval, respectively.

This ensures the test set remains aligned with the training distribution while incorporating astrophysical realism based on current observations.

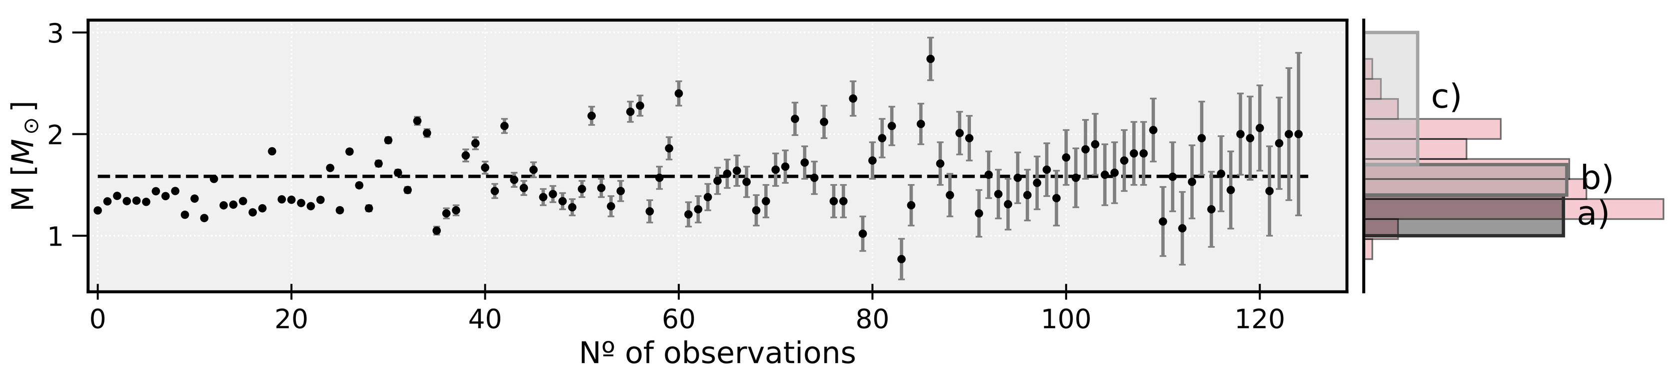

Neural Posterior Estimation

Neural Posterior Estimation (NPE) is a simulation-based inference method designed for amortized Bayesian inference. Instead of sampling from the posterior for each new observation—as in traditional approaches like MCMC—NPE trains a density estimator model, based on neural networks, to approximate the posterior distribution for any observation.

Given observed data $\boldsymbol{d}$ and parameters of interest $\boldsymbol{\theta}$, the goal is to learn a flexible approximation $q_{\boldsymbol{\phi}}(\boldsymbol{\theta}|\boldsymbol{d})$ to the true posterior $p(\boldsymbol{\theta}|\boldsymbol{d})$. Here, $\boldsymbol{\phi}$ denotes the trainable parameters of the neural network.

The training objective is to minimize the expected Kullback–Leibler divergence between the true posterior and the learned approximation:

$$L_{\text{NPE}} = \mathbb{E}_{\boldsymbol{d} \sim p(\boldsymbol{d})}\, \mathrm{D}_{\mathrm{KL}}\left[p(\boldsymbol{\theta}|\boldsymbol{d}) \,\|\, q_{\boldsymbol{\phi}}(\boldsymbol{\theta}|\boldsymbol{d})\right] \tag{1}$$

We implement $q_{\boldsymbol{\phi}}$ using *conditional normalizing flows* —a class of invertible neural networks that map a simple base distribution (e.g., Gaussian) into a complex target posterior. The transformation is defined via a sequence of invertible and differentiable functions:

$$q_{\boldsymbol{\phi}}(\boldsymbol{\theta}|\boldsymbol{d}) = p_z\left(f_{\boldsymbol{\phi}}^{-1}(\boldsymbol{\theta}; \boldsymbol{d})\right) \left|\det\left(\frac{\partial f_{\boldsymbol{\phi}}^{-1}}{\partial \boldsymbol{\theta}}\right)\right| \tag{2}$$

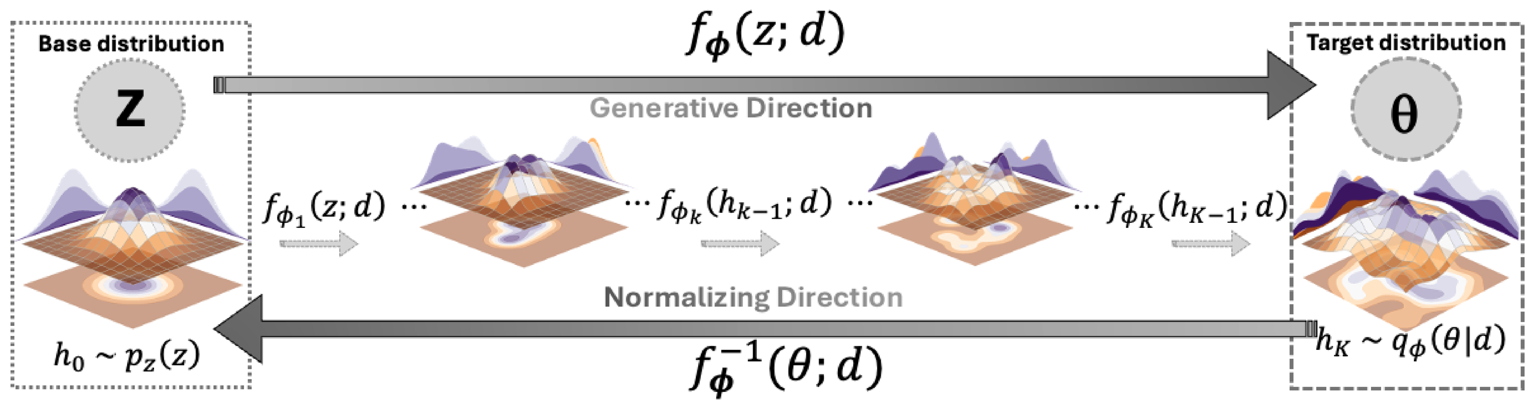

The total loss function used to train the flow-based model combines a likelihood-based term with a physics-informed regularization. It is defined as:

$$L = -\mathbb{E}_{\boldsymbol{\theta} \sim p(\boldsymbol{\theta})}\, \mathbb{E}_{\boldsymbol{d} \sim p(\boldsymbol{d}|\boldsymbol{\theta})} \left[\log p_z(f_{\boldsymbol{\phi}}^{-1}(\boldsymbol{\theta}; \boldsymbol{d})) + \log\left|\det \frac{\partial f_{\boldsymbol{\phi}}^{-1}}{\partial \boldsymbol{\theta}}\right|\right] + \lambda \cdot \sum_{i=1}^{19} \max\left(0,\, p(n_i) - p(n_{i+1})\right). \tag{3}$$

Output predictions

We illustrate predictions for two representative EoS from the test set in the left plot, the shaded band represents the 90% confidence interval (CI), and the dashed line shows the predicted median. The solid curve represent the true value and the dot denotes the maximum central baryonic density of the most massive neutron star for each EoS. In the right plot we are showing the coverage probability across the 20 pressure values. Both datasets are for dataset R_2.

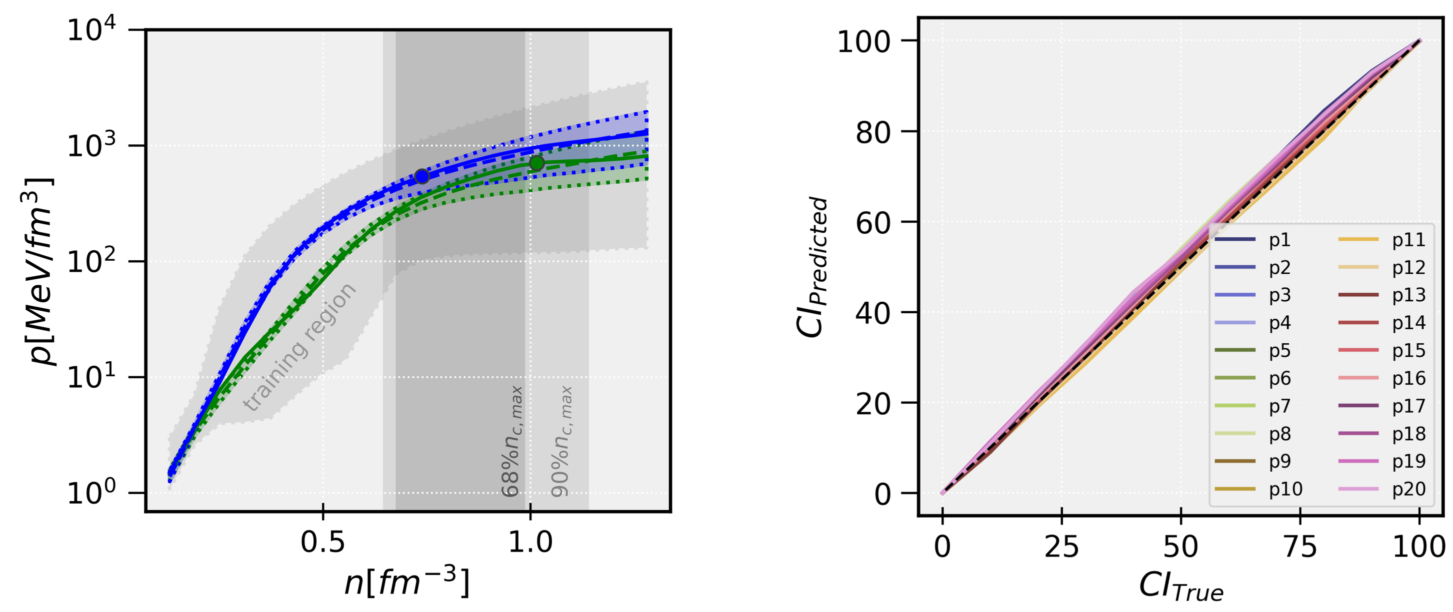

To better understand the uncertainty in predictions, we compute the normalized dispersion across the four datasets we created, represented in the table. This quantity is defined as the ratio of the 90% CI width to the median: $\frac{p_{90\%\text{CI}}}{\overline{p}}$ as a function of the normalized baryon density $n/n_{c,\text{max}}$, where $n_{c,\text{max}}$ is the maximum central density of the most massive stable star. The shaded region represents the 90% CI across all test models, and the dots with black edges are the mean. As expected, uncertainty increases with higher density (approaching $n/n_{c,\text{max}} = 1$), especially in models trained without observational noise, where the slope of the dispersion visibly steepens.

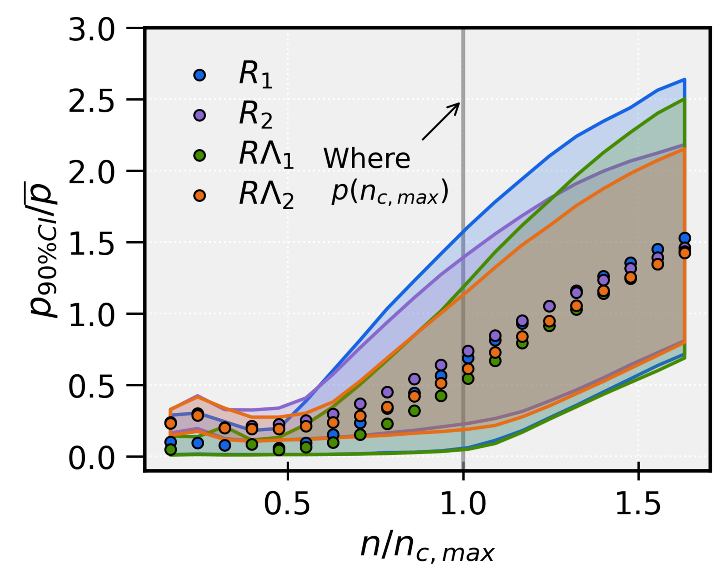

Dataset Configurations

Property / Set $\mathbb{X}$	R_1	R_2	$R\Lambda_1$	$R\Lambda_2$
Observational noise	✗	✔	✗	✔
Tidal deformability (Λ)	✗	✗	✔	✔

A check mark (✔) indicates presence, and a cross (✗) indicates absence.

Future work

- Apply the method to real observations from NICER and LIGO-Virgo-KAGRA.
- Extend the model to infer additional EoS-related quantities beyond pressure.
- Investigate more expressive normalizing flow architectures and hierarchical priors to better capture multimodal posteriors.
- Leverage HPC resources for large model ensembles and improved uncertainty quantification.

[1] Márcio Ferreira and Michał Bejger. Conditional variational autoencoder inference of neutron star equation of state from astrophysical observations. *Physical Review D*, 111(2):023035, 2025.

[2] Lívia S Rocha, Jorge E Horvath, Lucas M de Sá, Gustavo Y Chinen, Lucas G Barão, and Marcio GB de Avellar. Mass distribution and maximum mass of neutron stars: Effects of orbital inclination angle. *Universe*, 10(1):3, 2023.