Rapid neutron star equation of state inference with Normalising Flows

Jessica Irwin
University of Glasgow
INT workshop - Neutron Rich Matter on Heaven and Earth
June 2023

Inferring the NS EOS with GWs

- We use gravitational waves (GWs) to probe the neutron star (NS) equation of state (EOS) from a new perspective.

- GWs provide measurements of the gravitational mass and tidal deformability from a binary neutron star (BNS) merger.

- We can infer nuclear properties of the high density NS EOS and make statements on matter composition.

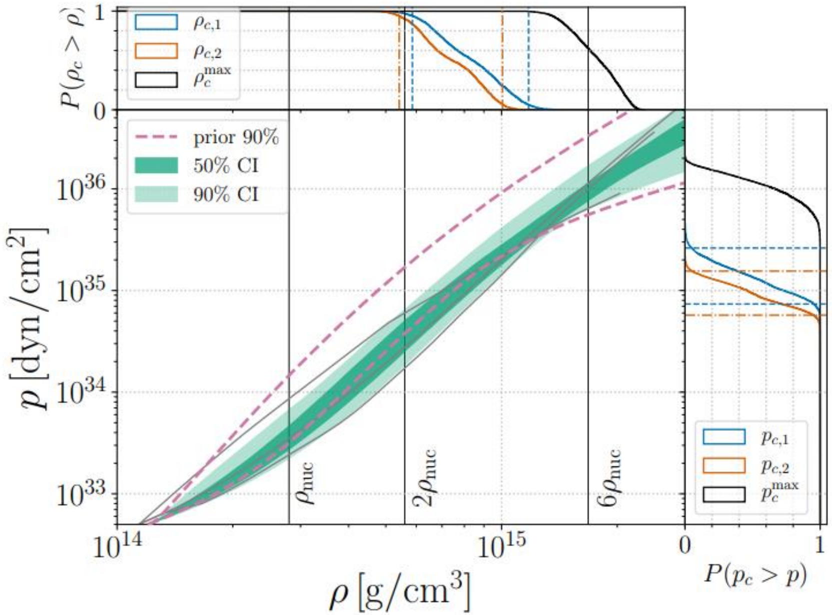

The LIGO Scientific Collaboration and The Virgo Collaboration (2017)

Neutron star tidal deformability

Tidal deformation that each star's gravitational field induces on its companion

$$\Lambda \equiv \frac{2}{3} k_2 C^{-5},$$

where k_2 is the love number and C is compactness

Astrophysical inference of Λ provides constraints of the NS EOS

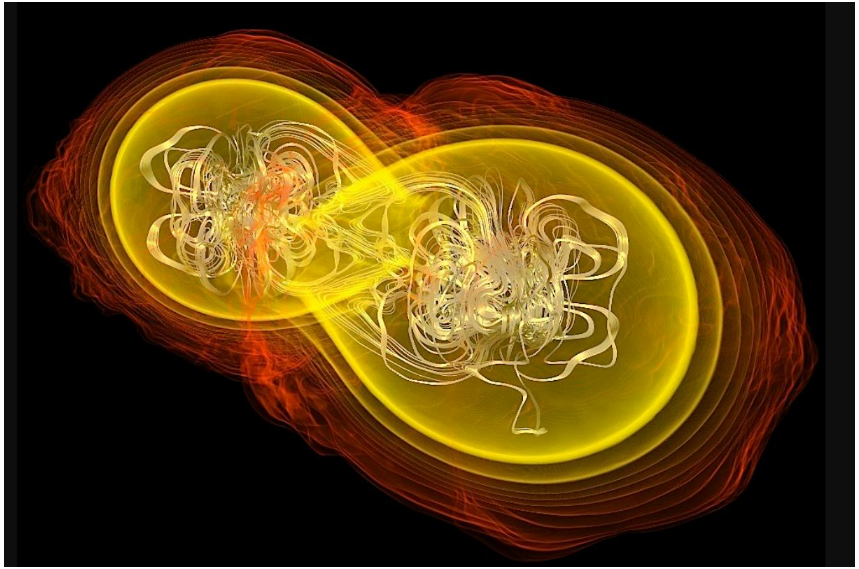

Measuring tidal deformability with GWs

Combined dimensionless quantity:

$$\tilde{\Lambda} = \frac{16}{13}\frac{(m_1 + 12m_2)m_1^4\Lambda_1 + (m_2 + 12m_1)m_2^4\Lambda_2}{(m_1 + m_2)^5}$$

Tidal effects significant at late stage of inspiral (merger)

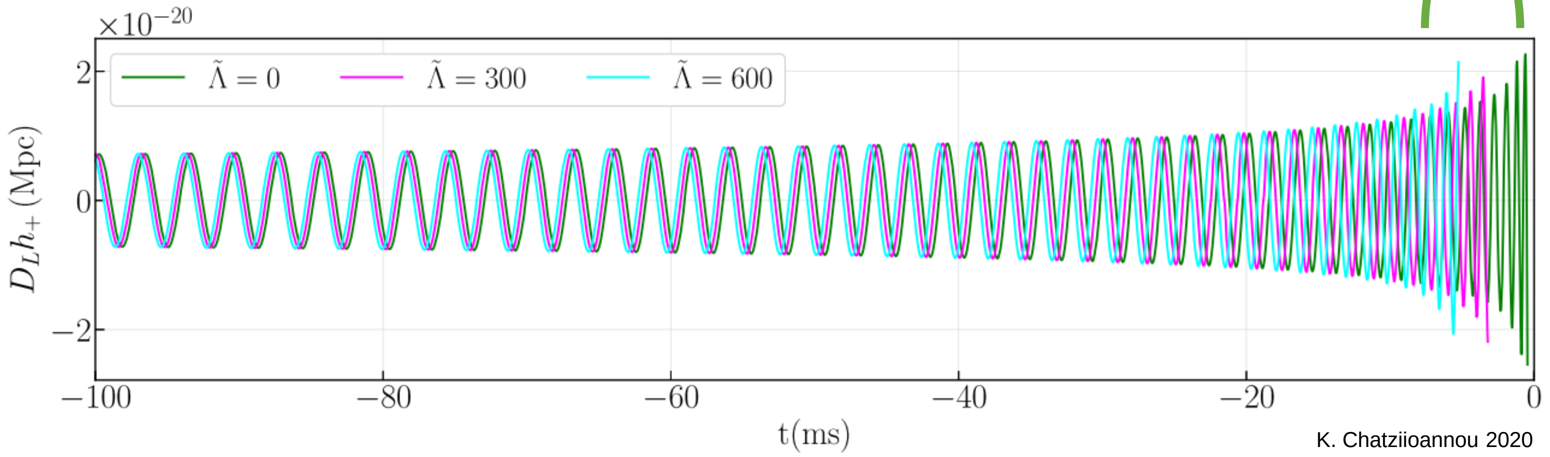

K. Chatziioannou 2020

Parameter estimation (PE)

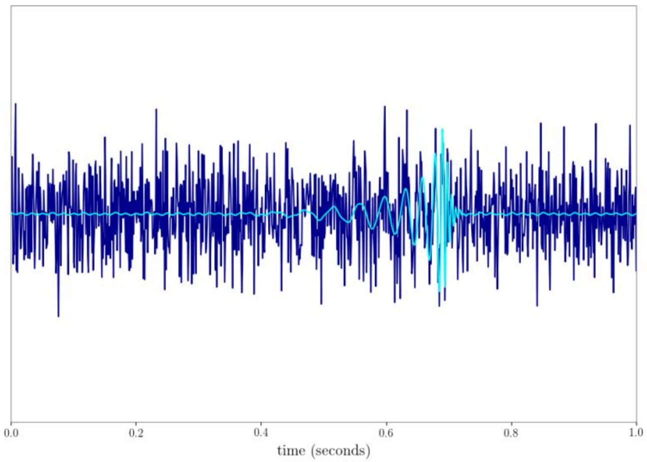

H. Gabbard et al. 2021

source properties:

$\mathcal{M}, q, d_L, \chi_1, \chi_2,$
$\delta, \alpha, \tilde{\Lambda}, \delta\tilde{\Lambda} \ldots$

Identification of signal → Matching to waveform model → Parameter estimation pipeline → Inference of source properties → **Equation of state inference**

Current uses of ML in GW data analysis

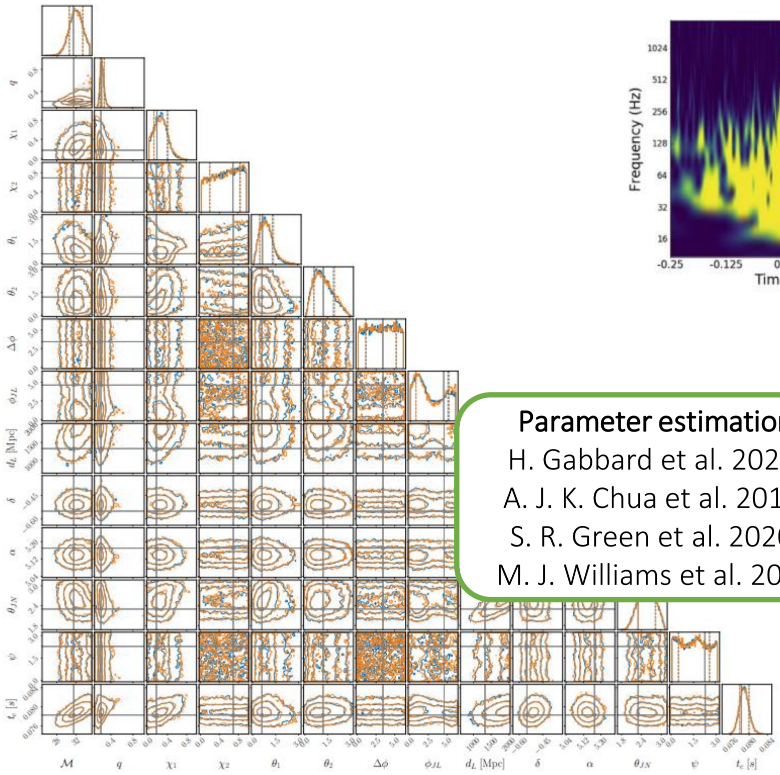

Parameter estimation
H. Gabbard et al. 2021,
A. J. K. Chua et al. 2019,
S. R. Green et al. 2020,
M. J. Williams et al. 2021

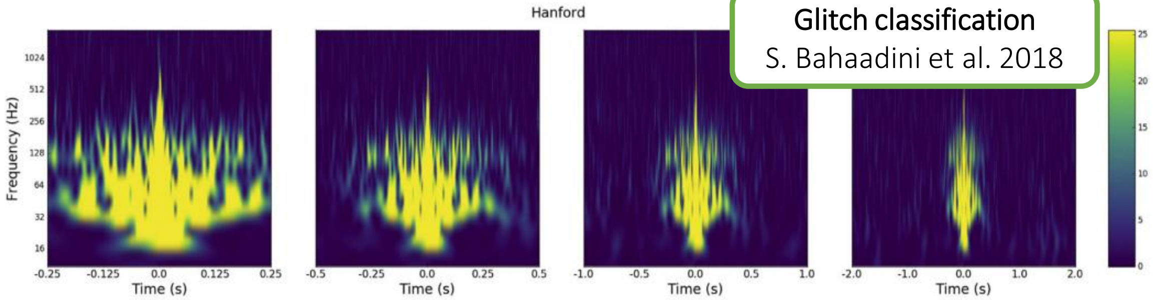

Glitch classification
S. Bahaadini et al. 2018

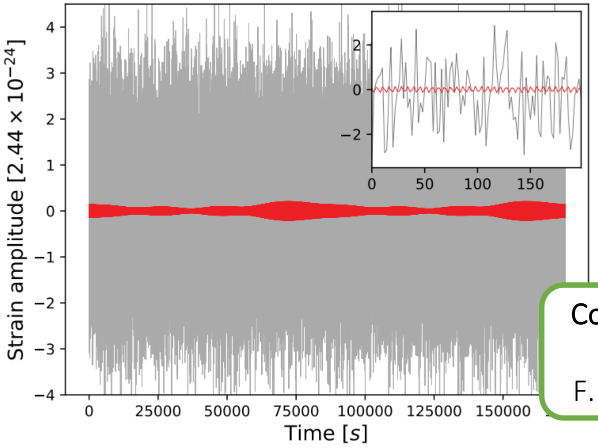

Continuous gravitational waves
F. Morawski et al. 2020

EOS inference with ASTREOS

- Machine learning tool that performs Bayesian inference of NS EOSs from PE samples of BNS events

- Return ensemble of non-parametric EOSs in pressure-density space alongside central and maximum density information

- Once trained, the Flow acts as a rapid functional generator of physically relevant EOSs

- Using a neural network has the advantage of being **intrinsically model independent** and **flexible.**

Normalising Flows

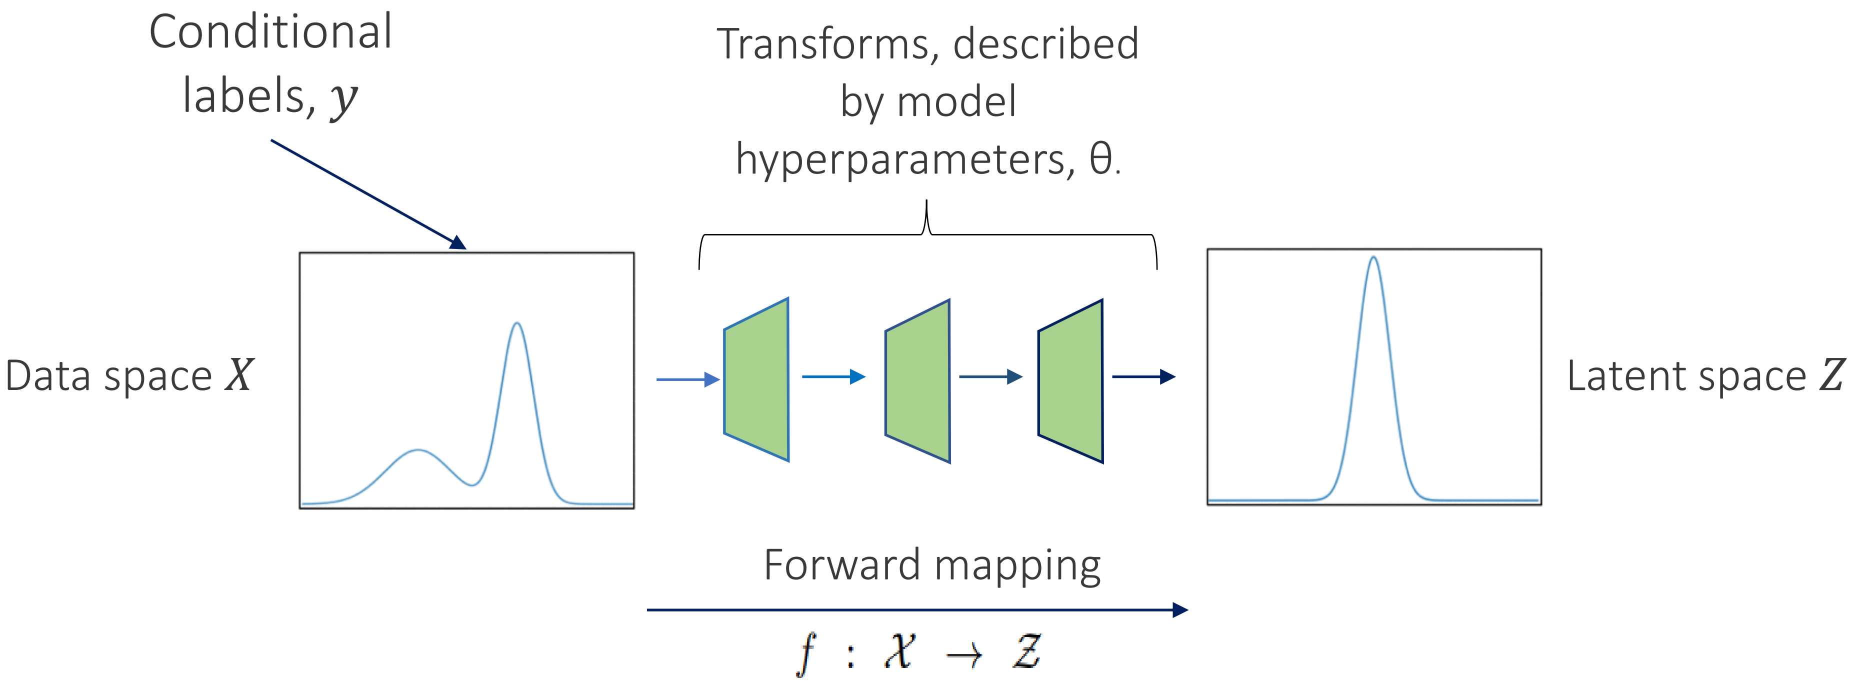

Images: I. Kobyzev, S. J. Prince, and M. A. Brubaker, 43, 3964–3979 (2021)

Normalising Flows

Conditional labels, y

Transforms, described by model hyperparameters, θ.

Data space X

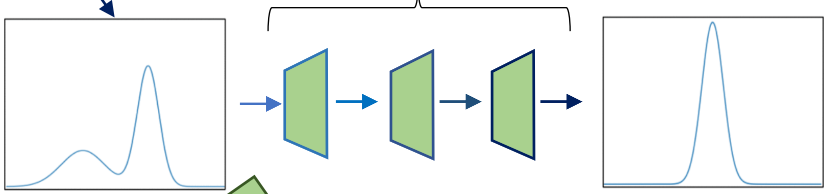

Latent space Z

Forward mapping

$f : \mathcal{X} \rightarrow \mathcal{Z}$

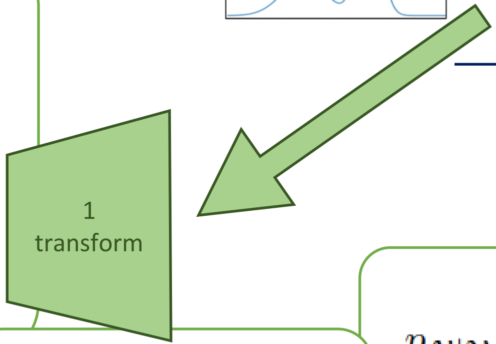

Global hyperparameters:
Epochs
Batch size
Learning rate
Learning rate scheduling
Batch normalisation

Model parameters θ:
Transforms
Blocks
Neurons

$$p_{\mathcal{X}|\mathcal{Y}}(x|y) = p_{\mathcal{Z}|\mathcal{Y}}(f(x|y)|y) \left|\det\left(\frac{\partial f(x|y)}{\partial x^T}\right)\right|$$

Goal of training:
Minimise the KL divergence between latent space Z and data space X

Physical conditions on equations of state

- Neural network has the advantage that it is **intrinsically model independent**
- Model information is stored in terms of weights and biases
- We generate 1x105 equation of state models to train the network
  - EOSs are piecewise polytropic
  - The lower density part of all EOSs is kept fixed to the SLy4 EOS
  - The higher density part of the EOS is variable and can produce different EOS profile from soft to stiff EOSs
  - The two regimes are matched at the cross-over density
  - We always ensure thermodynamic stability of the EOSs and check causality

EOS training data preparation

Select the equation of state

Each equation of state consists of energy density computed on a fixed grid of pressure.
Each equation of state consists of 105 points in the pressure-density space, truncated to retain high density information.

EOS training data preparation

Select the equation of state

Define component mass prior range and sample uniformly

Each EOS has a defined maximum mass. We define a uniform prior between $0.5M_{\odot}$ and the maximum mass allowed by this EOS. Component masses are then sampled ensuring that $m_1 \geq m_2$.

EOS training data preparation

EOS and maximum mass determines maximum energy density → component masses → component central energy densities → tidal deformability

Select the equation of state

↓

Define component mass prior range and sample uniformly

↓

Determine central densities and tidal deformability of components

EOS training data preparation

Select the equation of state

Define component mass prior range and sample uniformly

Determine central densities and tidal deformability of components

Result:

Training data: an EOS in energy density ρ on fixed grid of pressure P with corresponding central densities ρ_1 and ρ_2 and maximum allowed density ρ_{max} (auxiliary parameters)
Conditional data: an associated label $[m_1, m_2, \Lambda_1, \Lambda_2]$

Principal component analysis

- Use principal component analysis (PCA) to reduce the dimensionality of the training data space.
- Use 7 PCA components to represent the 105 data points in each EOS.
- Training data space then consists of 10 dimensions – 7 PCA components plus 3 auxiliary parameters (central densities of components and maximum allowed density)
- Total explained variance 99.975%

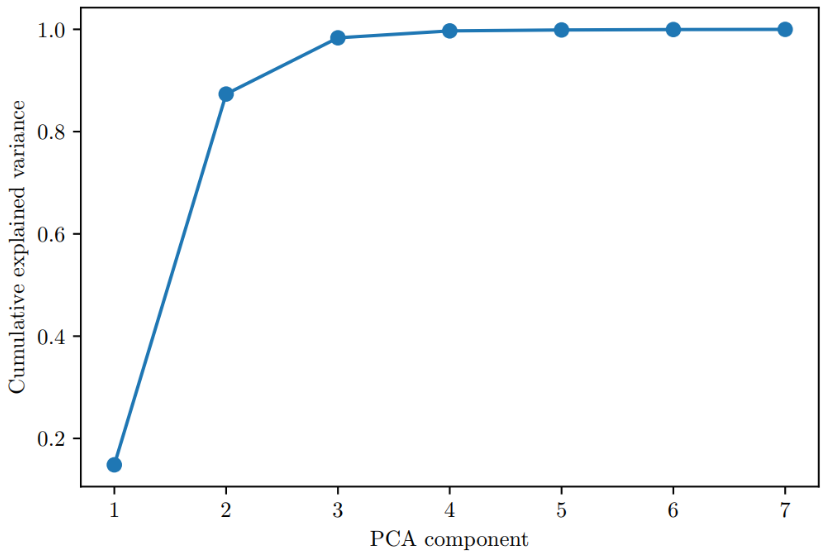

Training ASTREOS

Conditional labels, y:
$[m_1, m_2, \Lambda_1, \Lambda_2]$

3 transforms, each with 2 residual blocks consisting of 151 neurons

Data space, X:

EOS data (7 PCA) and 3 auxiliary parameters

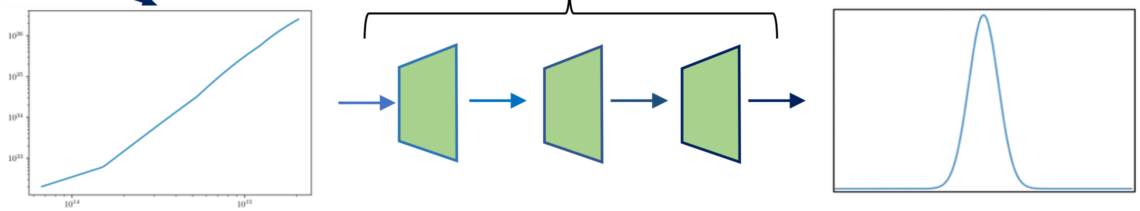

Latent space Z

Flow learns the forward mapping from EOS and auxiliary data conditional on NS parameters to multi-dimensional standard Gaussian

Training requires ~2 hours using a GPU (NVIDIA Tesla V100, ~2 GB memory footprint)

Using ASTREOS for EOS inference

3 transforms, each with 2 residual blocks consisting of 151 neurons

Conditional labels, y: $[m_1, m_2, \Lambda_1, \Lambda_2]$

Data space samples, x:

EOS + 3 auxiliary parameters

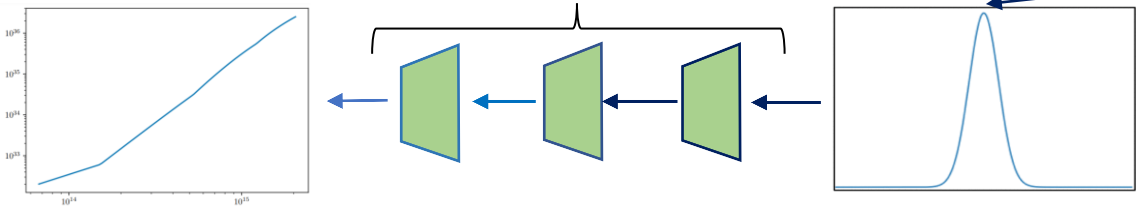

Latent space Z, learned by the Flow

Once trained, apply the inverse mapping from conditional label y and latent space location z to an EOS and corresponding auxiliary parameters

Can continually draw points from the latent space, natural variation in Flow output

Pre- and post-processing steps applied to ensure everything remains physically relevant

Results – GW170817 workflow

Inverse mapping

GW170817 EOS-insensitive posterior sample $[m_1, m_2, \Lambda_1, \Lambda_2]$ → Pre-processing → ASTREOS → Post-processing → EOS and auxiliary parameters

Pre-processing (21.7%)

- Pass EOS-insensitive PE samples from GW170817 to the Flow

Post-processing (5.3%, 19.9%)

- Generate 2500 EOSs + auxiliary parameters
- Reject EOSs based on processing steps
- Return EOS posterior distribution

Results – GW170817 workflow

Inverse mapping

GW170817 EOS-insensitive posterior sample $[m_1, m_2, \Lambda_1, \Lambda_2]$ → Pre-processing → ASTREOS → Post-processing → EOS and auxiliary parameters

Pre-processing (21.7%)

9.9%)

Pass EOS-insensitive PE samples from GW170817 to the Flow

2500 EOSs + auxiliary parameters

t EOSs based on processing steps

Return EOS posterior distribution

~0.1 seconds

Results – GW170817 EOS posterior

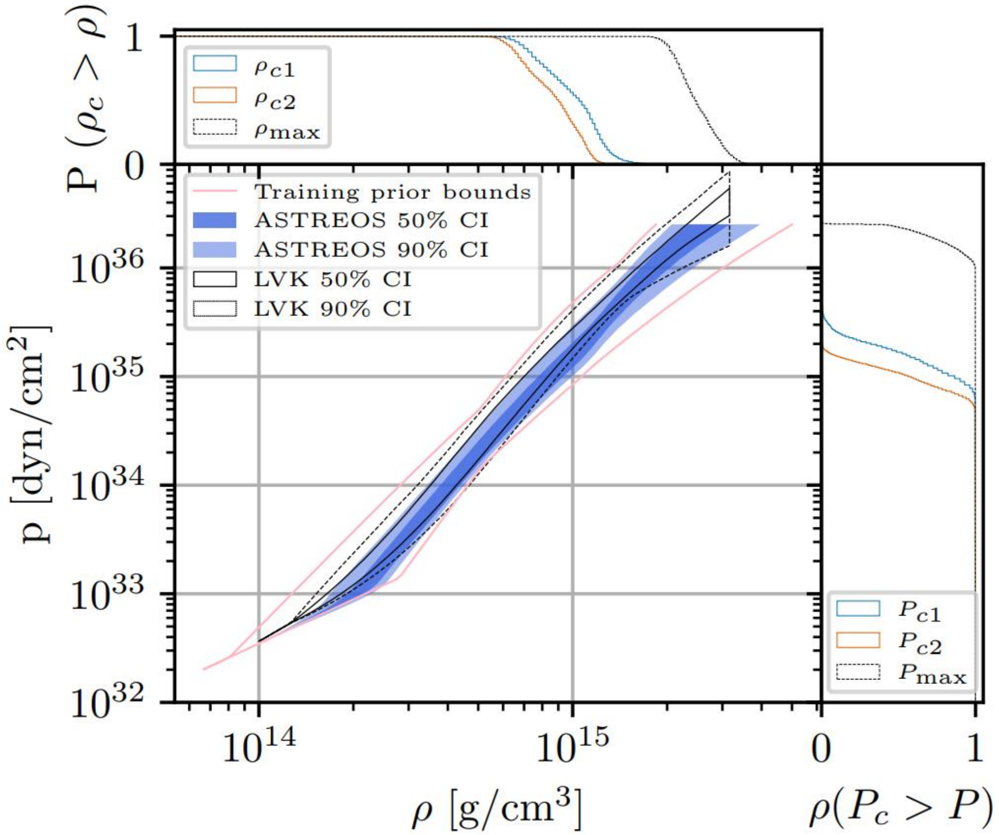

- Confidence intervals of 2500 EOS curves generated by passing EOS-insensitive posterior samples to the Flow
- ~0.1 seconds to generate 2500 EOS and auxiliary parameter samples
- ASTREOS demonstrates broad general agreement with LVK

Results & discussion

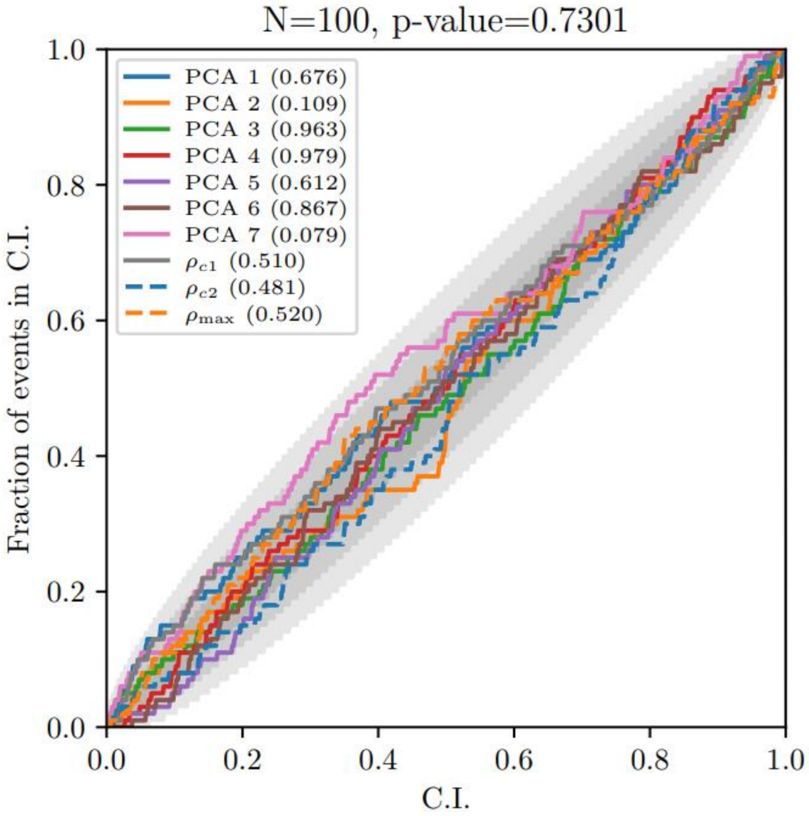

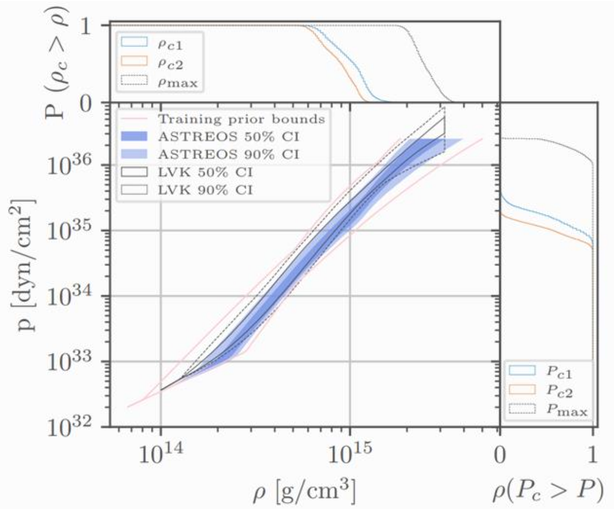

Pass 2000 samples of conditional label y to the flow to obtain 2000 EOS + auxiliary parameter samples

Isolated test demonstrates ASTREOS is statistically robust.

Conclusions and future work

- Flows can accurately and rapidly infer the neutron star equation of state
- Needs to be trained only once for repeatedly performing rapid inference for **all** possible future events
- Explicit model independent approach
- Easily modifiable for alternative conditional statements
- Complements existing literature and developments of ML in low latency GW science
- Easier to combine over multiple GW events and potential implications for BNS population inference

Based on work in prep:

Jordan McGinn, Arunava Mukherjee, Jessica Irwin, Christopher Messenger, Michael J. Williams, Ik Siong Heng

j.irwin.1@research.gla.ac.uk

Pre- and post-processing of data

Pre-processing

Place restrictions on conditional data y that can be passed to the Flow.

Construct a **convex hull** around the mass and tidal deformability prior space.

Do not allow posterior samples from a GW event outside the hull to be passed to the Flow.

Tidal deformability, central densities and maximum density are **standardised** to zero mean and unit variance

Post-processing

Gaussian mixture model (GMM) constructed to represent the 10-dimensional EOS + auxiliary parameter prior space.

Samples generated by the Flow that exist outside the GMM threshold are rejected.

In the $\rho - P$ space, EOSs that lie outside the **training prior bounds** are rejected, as these EOSs are assumed to be unphysical.

Conditional label prior space and processing result

21.7% of GW170817 EOS-insensitive posterior samples lie outside the convex hull constructed from our prior training data and are rejected

Blue is training prior space in $[m_1, m_2, \Lambda_1, \Lambda_2]$

Red are samples that lie outside the convex hull constructed around this training prior space

Green are samples that lie within the convex hull are then passed to the Flow as conditional labels

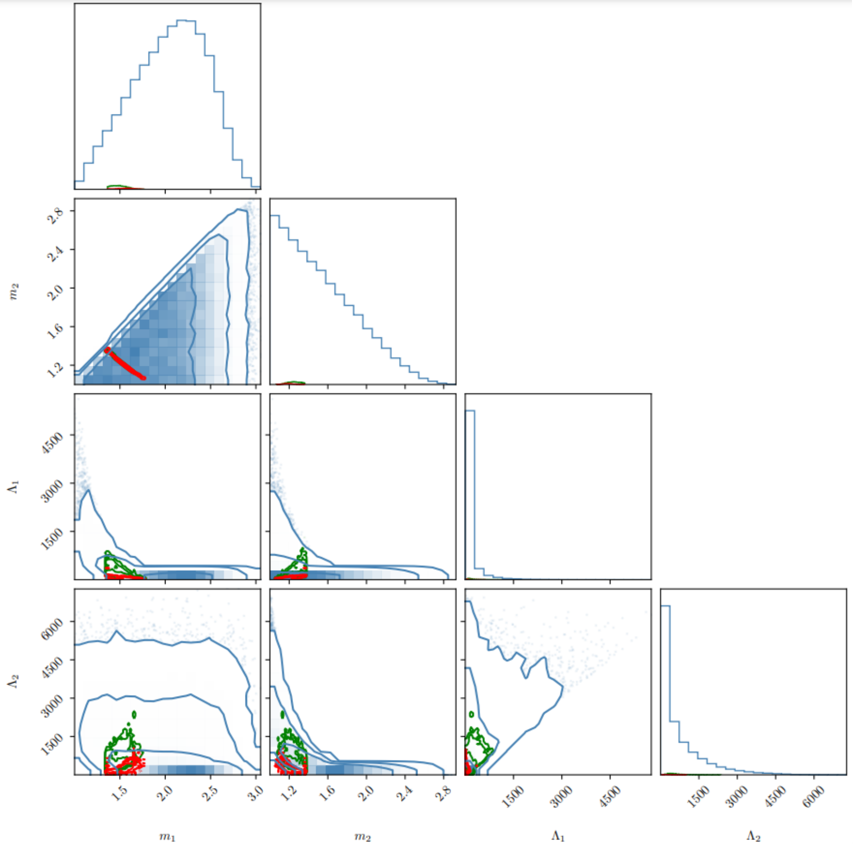

Conditional label prior space and processing result

Red are samples that lie outside the convex hull constructed around this training prior space.

Green are samples that lie within the convex hull and are then are passed to the Flow as conditional labels.

The boundary between the samples in and out of the hull is unclear as we are projecting a 4-dimensional data space.

The convex hull creates a 4-dimensional space which contains all training data samples in the mass-lambda space.

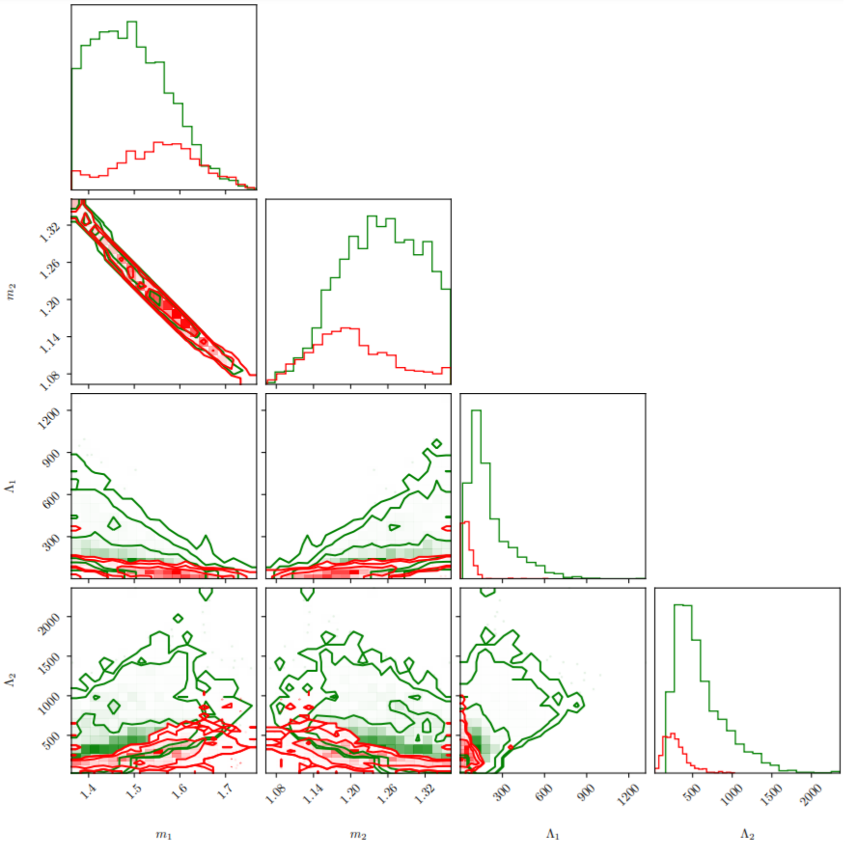

10-dimensional EOS prior space and GMM representation

5.3% of samples generated by the Flow (in the 10-dimensional PCA + auxiliary space) are rejected by the GMM threshold.

The 10-dimensional GMM is an approximation (see plot on right).

We find the outermost sample in the 10D PCA-reduced training data with respect to the GMM and set a threshold with respect to the probability of this sample belonging to the GMM.

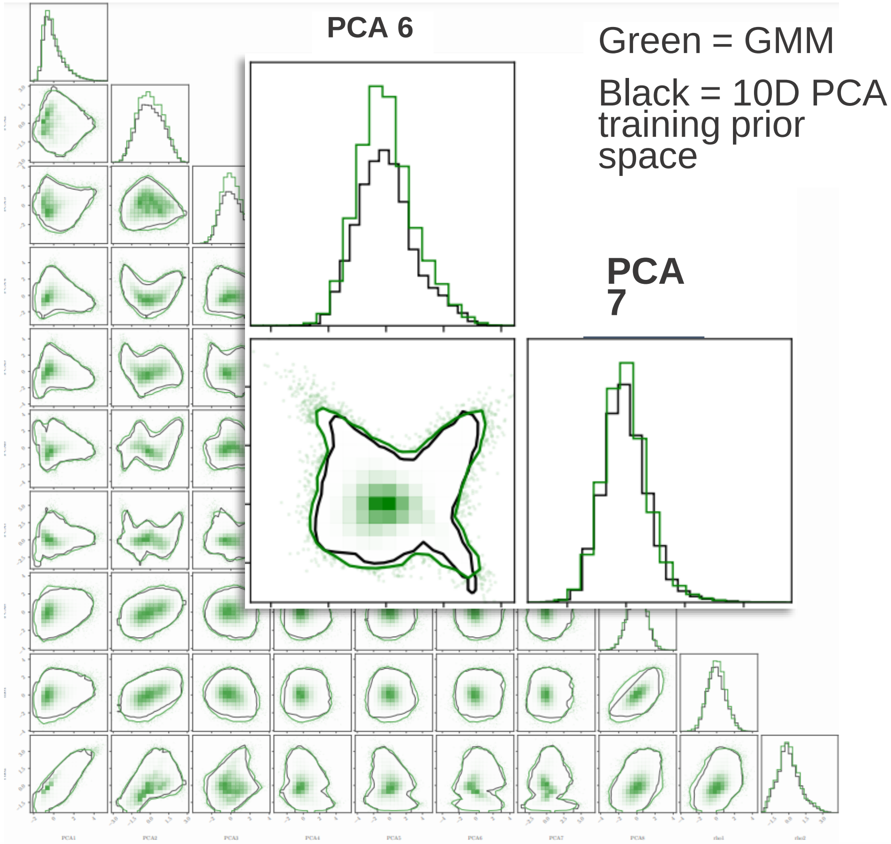

Rejecting EOSs with respect to training prior bounds

19.9% of remaining samples in the $\rho - P$ space are rejected due to lying outside the training prior boundaries.

The GMM rejection threshold was an approximation (see plot of GMM and prior space on slide 20) and so EOSs in the PCA-approximation may be accepted that exist near the probability threshold.

After projecting back into the pressure-density space, EOSs may be returned outside the training prior boundaries