Southampton Gravity Seminar
6 March 2025

ROYAL HOLLOWAY UNIVERSITY OF LONDON

Simulation-based inference for pulsar population synthesis

Dr. Vanessa Graber

in collaboration with Michele Ronchi, Celsa Pardo Araujo, and Nanda Rea

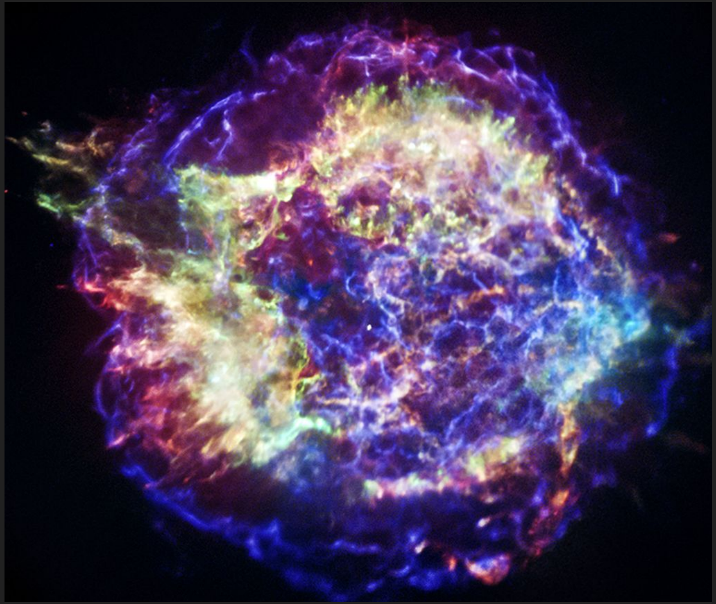

Cassiopeia A supernova remnant
(credit: NASA/CXC/SAO)

Outline

- Neutron stars
- Pulsar population synthesis
- Machine learning and sbi
- Inference results
- Summary

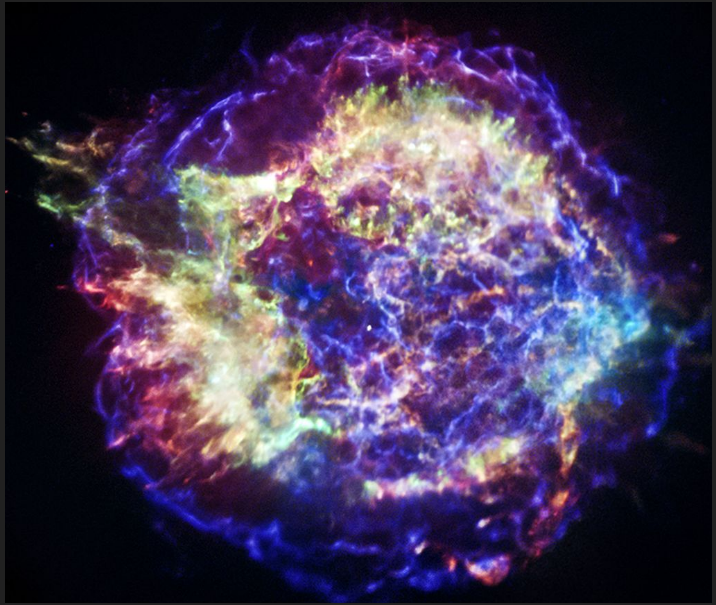

Cassiopeia A supernova remnant
(credit: NASA/CXC/SAO)

Neutron-star formation

- Neutron stars are one of three types of **compact remnants**, created during the **final stages of stellar evolution**.

- When a **massive star of 8 - 25 solar masses** runs out of fuel, it collapses under its own gravitational attraction and **explodes in a supernova**.

- During the collapse, **electron capture** processes ($p + e^- \rightarrow n + \nu_e$) produce (a lot of) neutrons.

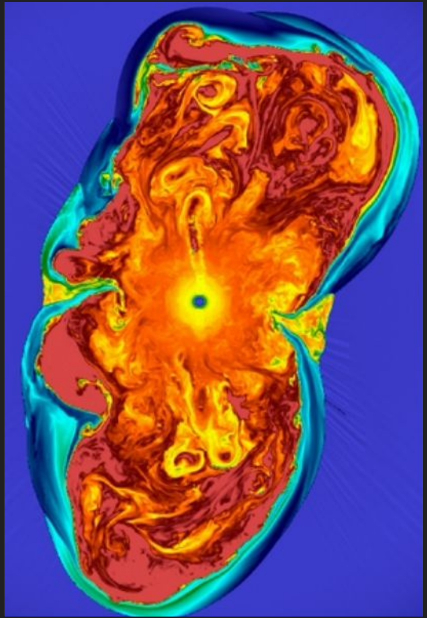

Snapshot of a 3D core-collapse supernova simulation (Mösta et al., 2014)

mass: 1.2 - 2.1 $M_{\odot}$

radius: 9 - 15 km

density: 10^{15} g/cm^3

Lighthouse radiation

- Neutron stars have **extreme magnetic fields** between 10^8 - 10^{15} G. For comparison, the Earth's magnetic field is 0.5 G.

- Because rotation and magnetic axes are misaligned, neutron stars emit radio beams **like a lighthouse**.

- These pulses can be observed with radio telescopes. This is how neutron stars were first detected and why we call them **pulsars**.

Sketch of the neutron-star exterior.

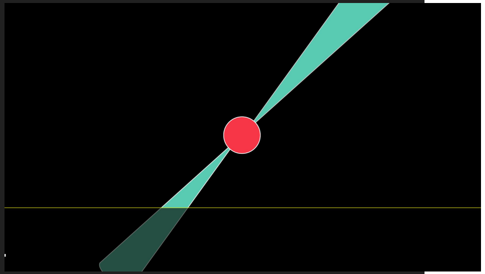

Credit: J. Christiansen

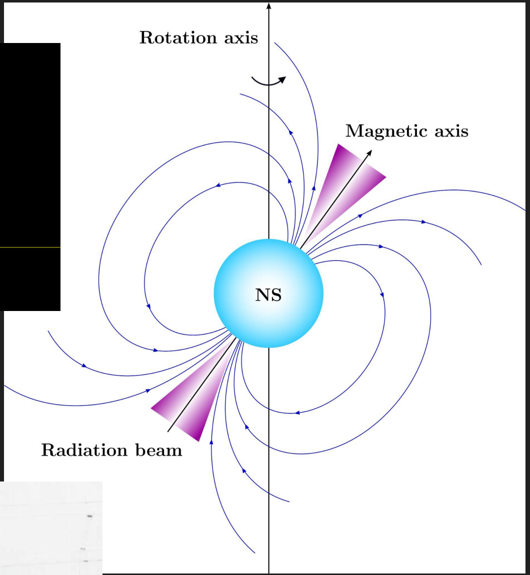

Dame Jocelyn Bell Burnell in front of her radio telescope in Cambridge, UK.

The neutron-star zoo

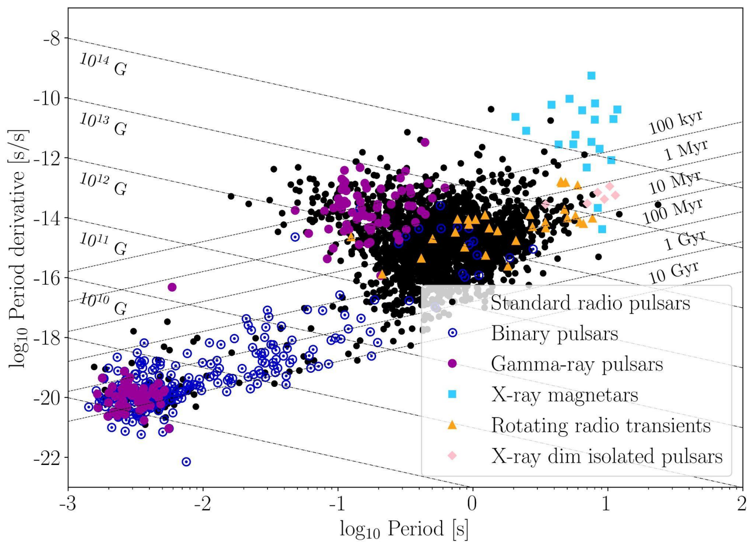

Period period-derivative plane for the pulsar population. Data taken from the ATNF Pulsar Catalogue (Manchester et al., 2005)

- Pulsars are **very precise clocks** and we time their pulses to **measure rotation periods P and derivatives $\dot{P}$**.
- We now observe neutron stars as pulsars **across** the **electromagnetic spectrum**.

~ **3,500 pulsars** are known to date

- Grouping neutron stars in the **$P\dot{P}$-plane** according to their observed properties serves as a diagnostic tool to **identify different neutron-star classes**.

The neutron-star zoo

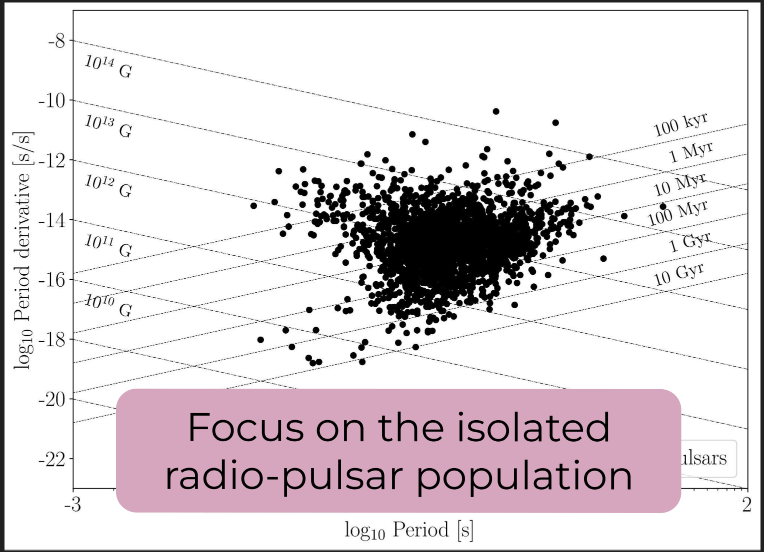

Period period-derivative plane for the pulsar population. Data taken from the ATNF Pulsar Catalogue (Manchester et al., 2005)

- Pulsars are **very precise clocks** and we time their pulses to **measure rotation periods P and derivatives $\dot{P}$**.
- We now observe neutron stars as pulsars **across** the **electromagnetic spectrum**.

~ **3,500 pulsars** are known to date

- Grouping neutron stars in the **$P\dot{P}$-plane** according to their observed properties serves as a diagnostic tool to **identify different neutron-star classes**.

Outline

- Neutron stars
- **Pulsar population synthesis**
- Machine learning and sbi
- Inference results
- Summary

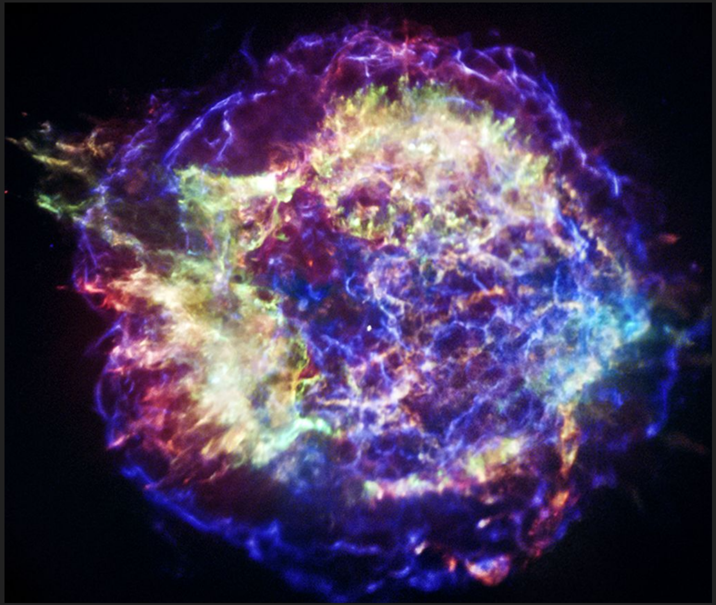

Cassiopeia A supernova remnant
(credit: NASA/CXC/SAO)

General idea

- We can estimate the **total number of neutron stars in our Galaxy**

CC supernova rate: ~ 2 per century	×	**Galaxy age:** ~ 13.6 billion years	=	**NS number:** ~ 2.8×10^8

- We only **detect** a very **small fraction** of all neutron stars. Population synthesis bridges this gap focusing on the full population of neutron stars (e.g. Faucher-Giguère & Kaspi 2006, Lorimer et al. 2006, Gullón et al. 2014, Cieślar et al. 2020):

model **birth properties** with Monte-Carlo approach → **evolve properties** forward in time → **apply filters** to mimic observational biases/limits → **compare** mock simulations **to observations** to constrain input

Goals

- Population synthesis allows us to **constrain the natal properties** of neutron stars and their **birth rates**.

- This is for example **relevant for**:
  - Massive star evolution
  - Gamma-ray bursts
  - Fast-radio bursts
  - Peculiar supernovae

Estimated Galactic core-collapse supernova rate and birth rates for different pulsar classes (Keane & Kramer, 2008).

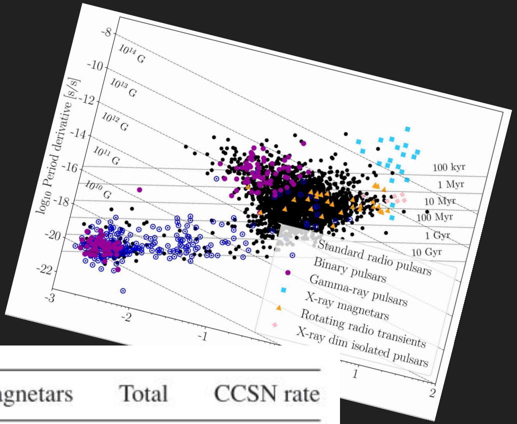

PSRs	RRATs	XDINSs	Magnetars	Total	CCSN rate
2.8 ± 0.5	$5.6^{+4.3}_{-3.3}$	2.1 ± 1.0	$0.3^{+1.2}_{-0.2}$	$10.8^{+7.0}_{-5.0}$	1.9 ± 1.1
1.4 ± 0.2	$2.8^{+1.6}_{-1.6}$	2.1 ± 1.0	$0.3^{+1.2}_{-0.2}$	$6.6^{+4.0}_{-3.0}$	1.9 ± 1.1
1.1 ± 0.2	$2.2^{+1.7}_{-1.3}$	2.1 ± 1.0	$0.3^{+1.2}_{-0.2}$	$5.7^{+4.1}_{-2.7}$	1.9 ± 1.1
1.6 ± 0.3	$3.2^{+2.5}_{-1.9}$	2.1 ± 1.0	$0.3^{+1.2}_{-0.2}$	$7.2^{+5.0}_{-3.4}$	1.9 ± 1.1
1.1 ± 0.2	$2.2^{+1.7}_{-1.3}$	2.1 ± 1.0	$0.3^{+1.2}_{-0.2}$	$5.7^{+4.1}_{-2.7}$	1.9 ± 1.1

- We can also learn about **evolutionary links between different neutron-star classes** (e.g., Viganó et al., 2013). This is important because estimates for the **Galactic core-collapse supernova rate** are **insufficient** for to explain the independent formation of different classes of pulsars (Keane & Kramer, 2008).

Dynamical evolution I

- **Neutron stars are born in star-forming regions**, i.e., in the Galactic disk along the Milky Way's spiral arms, **and receive kicks** during the supernova explosions.

- We make the following assumptions:
  - Spiral-arm model (Yao et al., 2017) plus rigid rotation with T = 250 Myr
  - **Exponential disk model** with scale height h_c (Wainscoat et al., 1992)
  - Single-component **Maxwell kick-velocity distribution** with dispersion σ_k (Hobbs et al., 2005)
  - Galactic potential (Marchetti et al., 2019)

Artistic illustration of the Milky Way (credit: NASA JPL)

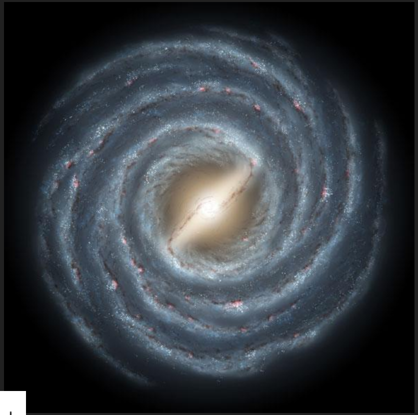

$$\mathcal{P}(z) = \frac{1}{h_c}\, e^{-\frac{|z|}{h_c}}$$

$$\mathcal{P}(v_\mathrm{k}) = \sqrt{\frac{2}{\pi}}\frac{v_\mathrm{k}^2}{\sigma_\mathrm{k}^3}\, e^{-\frac{v_\mathrm{k}^2}{2\sigma_\mathrm{k}^2}}$$

For Monte-Carlo approach, we **vary two uncertain parameters h_c and σ_k.**

Vanessa.Graber@rhul.ac.uk

Dynamical evolution II

- For our Galactic model Φ_{MW}, we evolve the stars' position & velocity by **solving Newtonian equations of motion** in cylindrical galactocentric coordinates:

$$\ddot{\vec{r}} = -\vec{\nabla}\Phi_{MW}$$

Top view

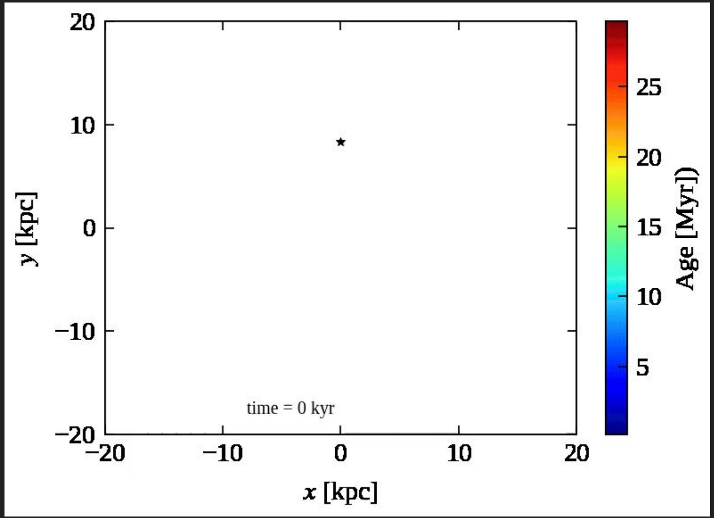

Side view

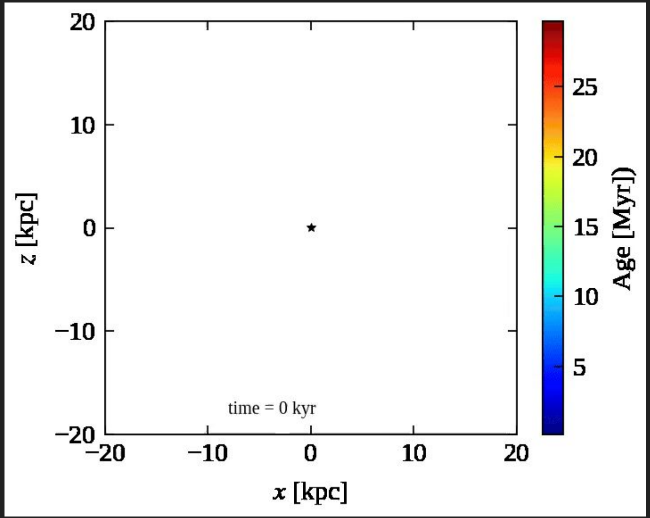

Galactic evolution tracks for h_c = 0.18 kpc, σ = 265 km/s.

Magneto-rotational evolution I

- The neutron-star magnetosphere exerts a **torque onto the star**. This causes **spin-down** and **alignment of the magnetic and rotation axes**.
- Neutron star **magnetic fields decay** due to the Hall effect and Ohmic dissipation in the outer stellar layer (crust) (e.g., Viganó et al., 2013 & 2021).
- We make the following assumptions:
  - **Initial periods** follow a log-normal with μ_P and σ_P (Igoshev et al., 2022)
  - **Initial fields** follow a log-normal with μ_B and σ_B (Gullón et al., 2014)
  - Above $\tau \sim 10^6$ yr, **field decay** follows a power-law with $B(t) \sim B_0\,(1 + t/\tau)^\alpha$.

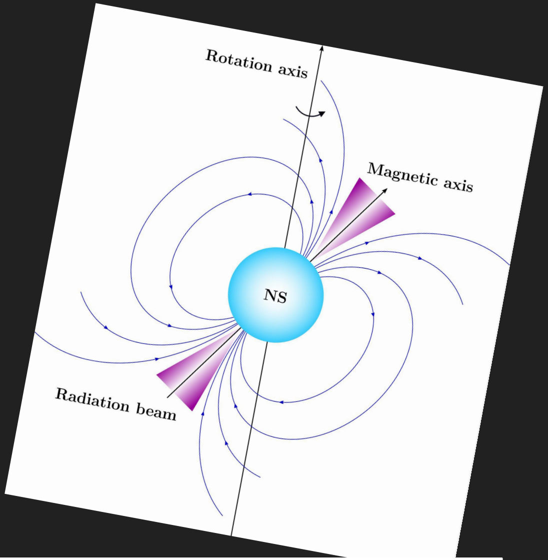

$$\mathcal{P}(\log P_0) = \frac{1}{\sqrt{2\pi}\sigma_P}\exp\left(-\frac{[\log P_0 - \mu_P]^2}{2\sigma_P^2}\right)$$

Here, we **vary** the five uncertain parameters **μ_P, μ_B, σ_P, σ_B** and **α**.

Vanessa.Graber@rhul.ac.uk

Magneto-rotational evolution II

- To model the magneto-rotational evolution, we numerically **solve two coupled ordinary differential equations** for the period and the misalignment angle (Aguilera et al., 2008; Philippov et al. 2014).

- We use results from **2D magneto-thermal simulations** to determine the evolution of the magnetic field below 10^6 yr (Viganò et al. 2021).

- This allows us to follow the stars' P and $\dot{P}$ evolution in the $P\dot{P}$-plane.

$P\dot{P}$ evolution tracks for $\mu_P = -0.6$, $\sigma_P = 0.3$, $\mu_B = 13.25$ and $\sigma_B = 0.75$.

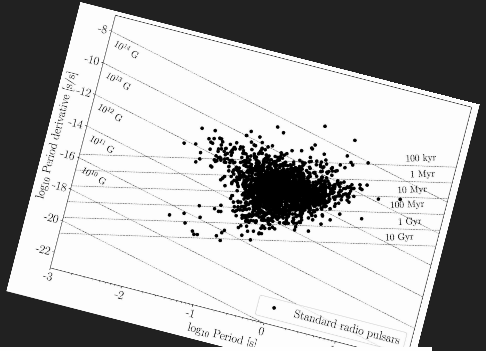

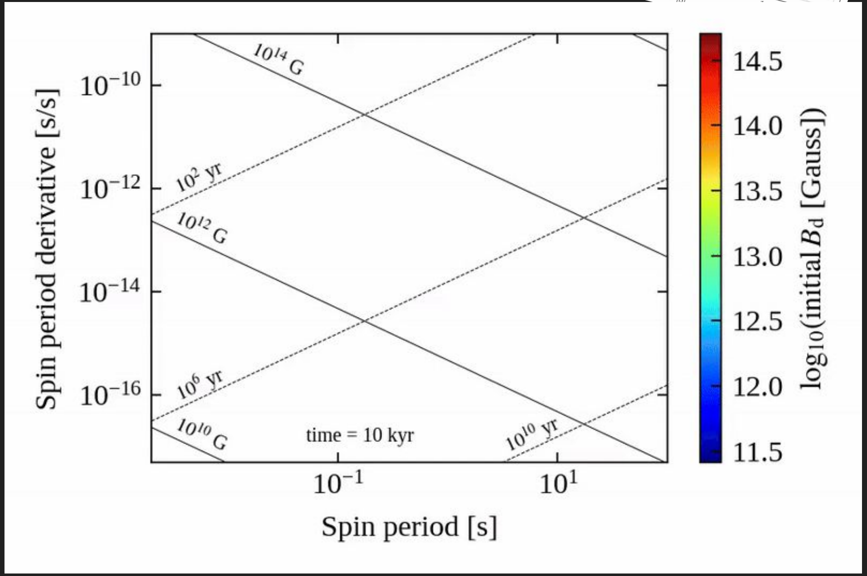

Radio emission and detection

- The stars' **rotational energy E_{rot}** is converted into coherent radio emission. We assume that the corresponding **radio luminosity L_{radio}** is proportional to the loss of E_{rot} (Faucher- Giguère & Kaspi, 2006; Gullón et al., 2014). L_0 is taken from observations.

$$L_{\rm radio} = L_0 \left(\frac{\dot{P}}{P^3} \right)^{1/2} \propto \dot{E}_{\rm dot}^{1/2}$$

- As **emission is beamed**, ~ 90% of pulsars do not point towards us. For those intercepting our line of sight, compute **radio flux S_{radio}** & **pulse width W**.

$$S_{\rm radio} = \frac{L_{\rm radio}}{\Omega_{\rm beam} d^2}$$

A pulsar counts as detected, if it **exceeds the sensitivity threshold** for a survey recorded with a specific radio telescope.

Three pulsar surveys

- We compare our simulated populations with three surveys from Murriyang (the Parkes Radio Telescope):
    - **Parkes Multibeam Pulsar Survey** (PMPS): *1,009 PSRs*
    - **Swinburne Parkes Multibeam Pulsar Survey** (SMPS): *218 PSRs*
    - **High Time Resolution Universe Survey** (HTRU): *1,023 PSRs*

Can we constrain birth properties by looking at a current snapshot of the pulsar population?

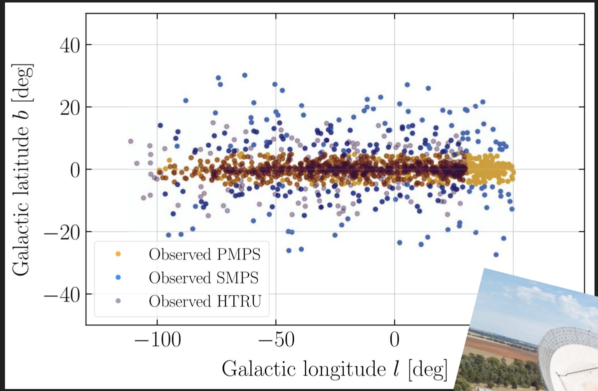

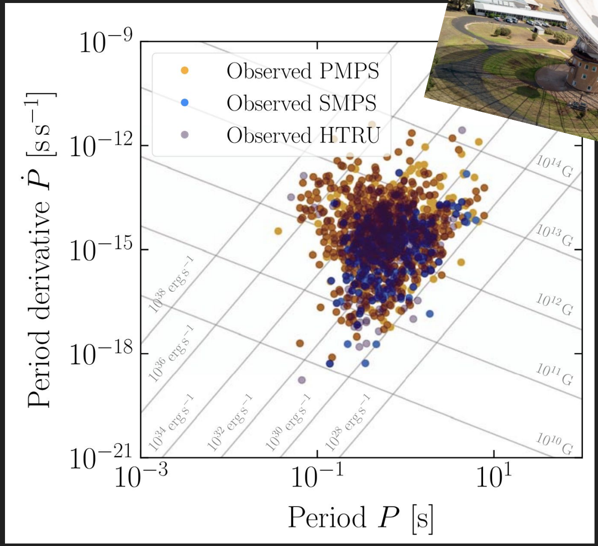

Outline

- Neutron stars
- Pulsar population synthesis
- **Machine learning and sbi**
- Inference results
- Summary

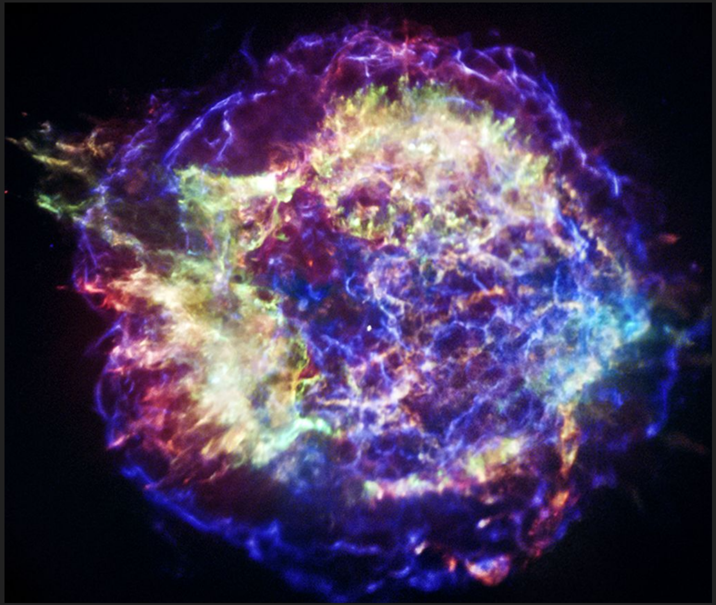

Cassiopeia A supernova remnant
(credit: NASA/CXC/SAO)

Comparing models and data

- Comparing observations to models and **constraining regions of the parameter space** that are **most probable given the data** is fundamental to many fields of science.

- Pulsar population synthesis is complex and has **many free parameters**. To compare synthetic simulations with observations, people have
  - Randomly sampled and then optimised 'by eye' (e.g., Gonthier et al., 2007)
  - Compared distributions of individual parameters using χ^2- and KS-tests (e.g., Narayan & Ostriker, 1990; Faucher-Giguère & Kaspi, 2006)
  - Used annealing methods for optimisation (Gullón et al., 2014)
  - Performed Bayesian inference for simplified models (Cieślar et al., 2020)

These methods do not scale well and are **difficult to use** with the **multi-dimensional data** produced in population synthesis.

Vanessa.Graber@rhul.ac.uk

Statistical inference

- Instead of **deducing point estimates**, we often do not require exact estimates but want to obtain **knowledge of probable regions**.

- This is where **Bayesian inference** comes in: based on some prior knowledge π (θ), a stochastic model and some observation x, we want to infer the most likely distribution P(θ|x) for our model parameters θ given the data x. This is **encoded in Bayes' Theorem**:

$$\underbrace{\mathcal{P}(\theta|x)}_{\text{posterior}} = \frac{\overbrace{\mathcal{P}(\theta)}^{\text{prior } \pi} \; \overbrace{\mathcal{P}(x|\theta)}^{\text{likelihood } \mathcal{L}}}{\underbrace{\mathcal{P}(x)}_{\text{evidence}}} = \frac{\mathcal{P}(\theta) \int \mathcal{P}(x, z|\theta) \mathrm{d}z}{\int \mathcal{P}(x|\theta')\mathcal{P}(\theta')\mathrm{d}\theta'}$$

For complex simulators, the **likelihood is defined implicitly and often intractable**. This is overcome with **simulation-based** (likelihood-free) **inference** (see e.g. Cranmer et al., 2020).

Vanessa.Graber@rhul.ac.uk

Simulation-based inference I

- To perform **Bayesian inference for any kind of (stochastic) forward model** (e.g. those specified by simulators), we use the following approach:

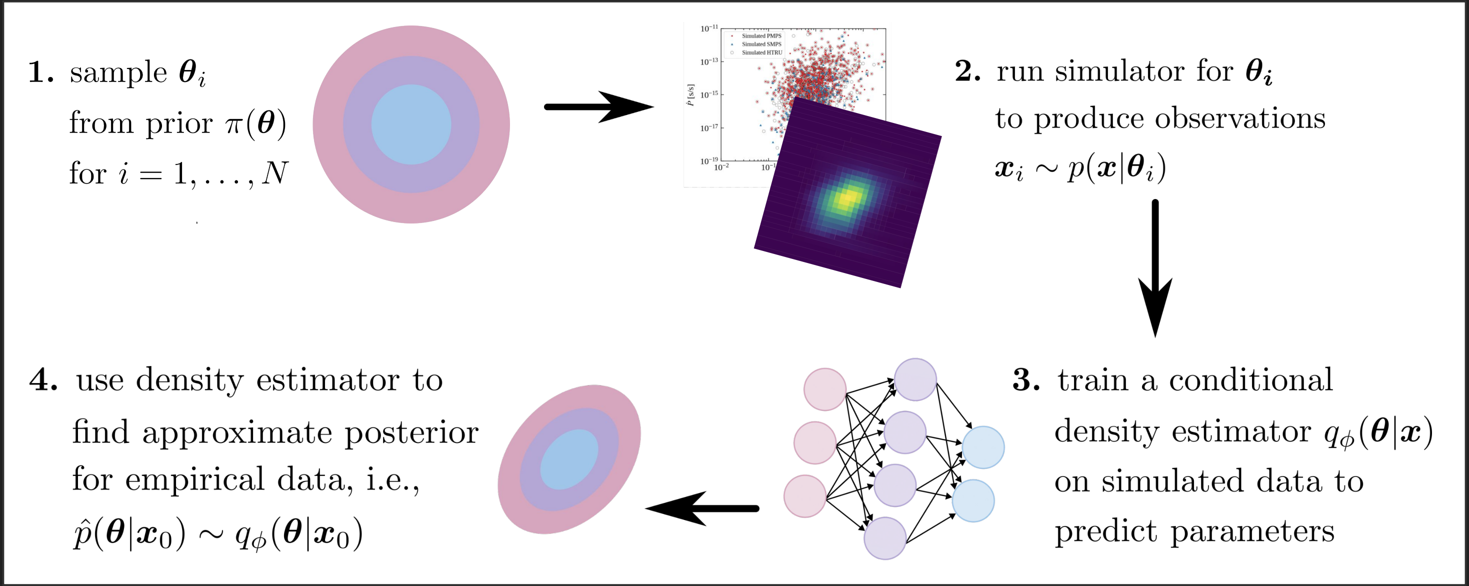

Simulation-based inference II

- Different approaches (all relying on deep learning) exist to **learn a probabilistic association** between the simulated data and the underlying parameters. These algorithms essentially focus on different pieces of Bayes' theorem:
  - Neural Posterior Estimation (NPE) (e.g., Papamakarios & Murray, 2016)
  - Neural Likelihood Estimation (NLE) (e.g., Papamakarios et al., 2019)
  - Neural Ratio Estimation (NRE) (e.g., Hermans et al., 2020; Delaunoy et al., 2022)

We focus on NPE. This allows us to **directly learn the posterior distribution**. In contrast, NLE and NRE need an extra (potentially time consuming) MCMC sampling step to construct a posterior.

- All methods exist in **sequential form** (SNPE, SNLE, SNRE), **which adds a fifth step to workflow**. Instead of sampling from the prior, we adaptively generate simulations from the posterior. This **typically requires fewer simulations**.

Simulation-based inference I

- To perform **Bayesian inference for any kind of (stochastic) forward model** (e.g. those specified by simulators), we use the following approach:

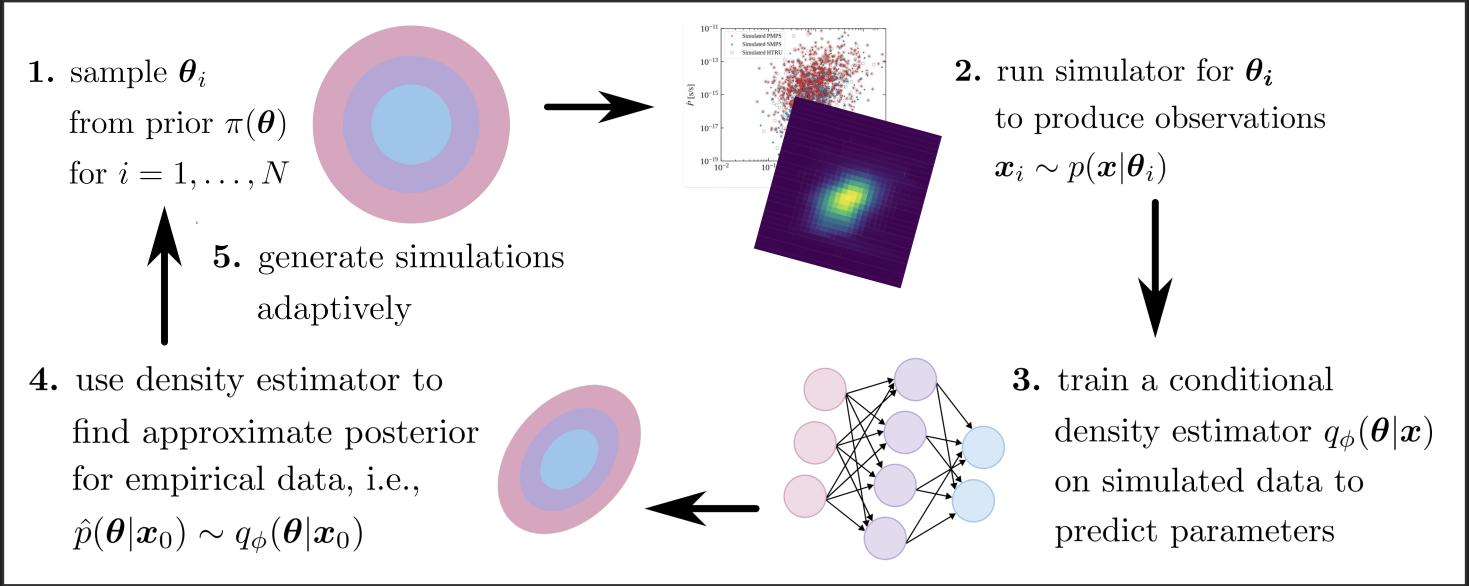

Outline

- Neutron stars
- Pulsar population synthesis
- Machine learning and sbi
- **Inference results**
- Summary

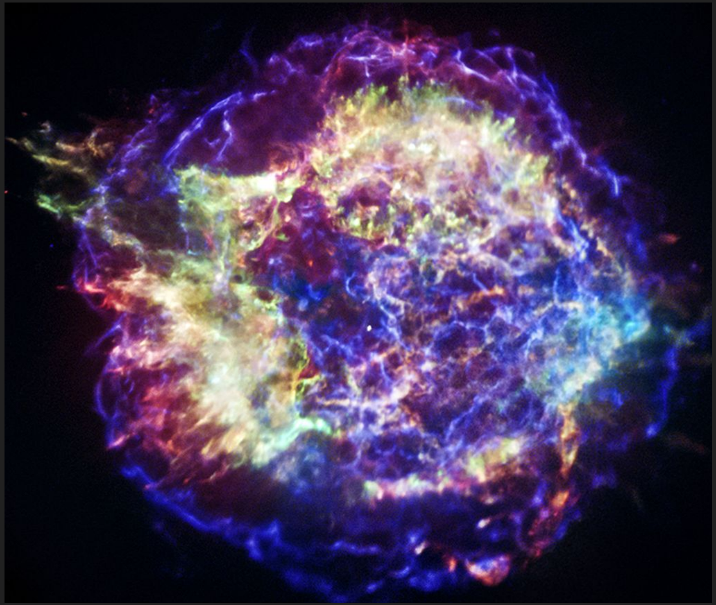

Cassiopeia A supernova remnant
(credit: NASA/CXC/SAO)

Workflow Graber et al. (2024)

Varying the five parameters $\boldsymbol{\mu_P}$, $\boldsymbol{\mu_B}$, $\boldsymbol{\sigma_P}$, $\boldsymbol{\sigma_B}$ and $\boldsymbol{\alpha}$, we **simulate 360,000 synthetic pulsar populations** over 6 weeks.

- With our complex population synthesis simulator, we fix the dynamics to a fiducial model and **focus on the magneto-rotational evolution**.
- From our simulated populations, we **generate summary statistics**: density maps for three surveys in the $P\dot{P}$-plane.
- Build on **PyTorch package sbi** (Tejero-Cantero et al., 2020):

$$
\begin{aligned}
&\mu_{\log B} \in \mathcal{U}(12, 14),\\
&\sigma_{\log B} \in \mathcal{U}(0.1, 1),\\
&\mu_{\log P} \in \mathcal{U}(-1.5, -0.3),\\
&\sigma_{\log P} \in \mathcal{U}(0.1, 1),\\
&a_{\mathrm{late}} \in \mathcal{U}(-3, -0.5).
\end{aligned}
$$

Workflow Graber et al. (2024)

- With our complex population synthesis simulator, we fix the dynamics to a fiducial model and **focus on the magneto-rotational evolution**.
- From our simulated populations, we **generate summary statistics**: density maps for three surveys in the $P\dot{P}$-plane.
- Build on **PyTorch package sbi** (Tejero-Cantero et al., 2020):

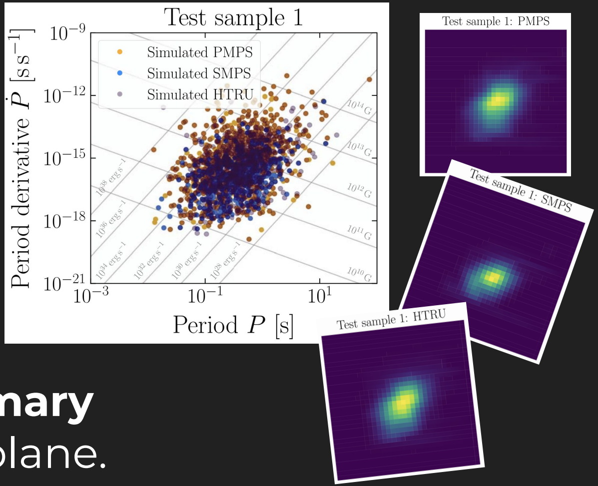

$$\mu_{\log B} \in \mathcal{U}(12,\ 14),$$
$$\sigma_{\log B} \in \mathcal{U}(0.1,\ 1),$$
$$\mu_{\log P} \in \mathcal{U}(-1.5,\ -0.3),$$
$$\sigma_{\log P} \in \mathcal{U}(0.1,\ 1),$$
$$a_{\rm late} \in \mathcal{U}(-3,\ -0.5).$$

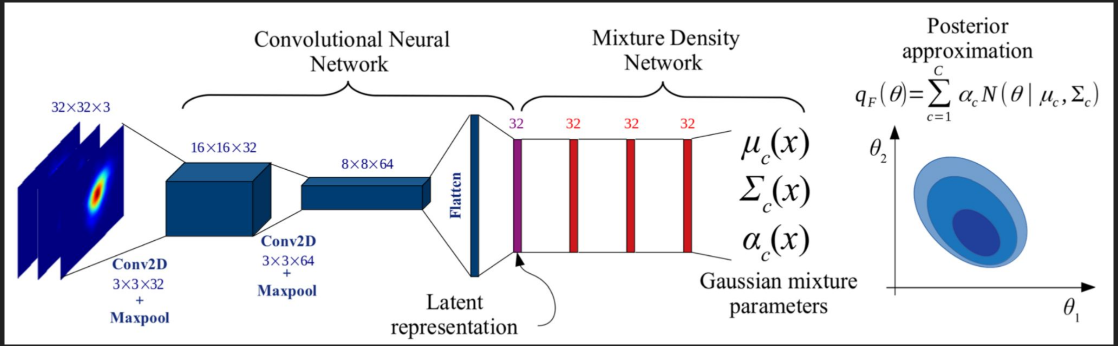

Results I

- Inferring on test samples with known ground truths, we recover narrow and well-defined posteriors for all parameters.

- We varied the hyperparameters of our DL approach to test robustness, and find very similar training behaviour and optimisation losses for 19 inference experiments.

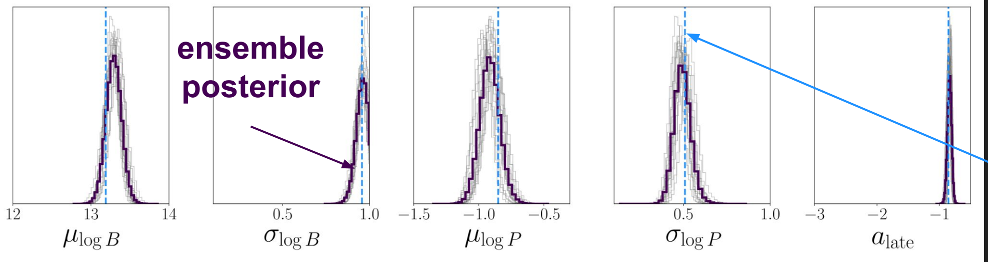

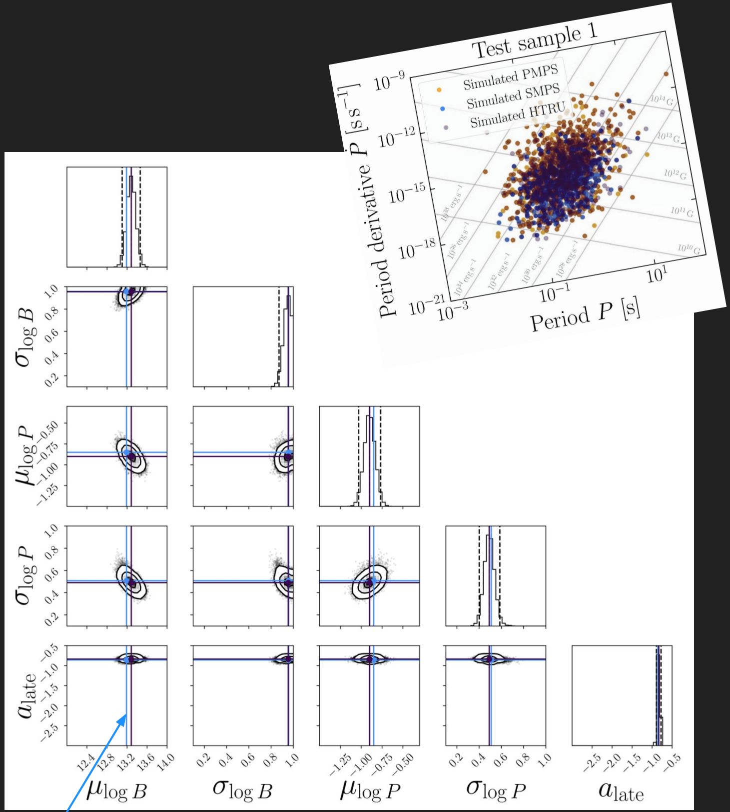

ground truths

Results II

- We then use the **ensemble posterior** to infer on the observed population:
  - Initial B-field parameters are narrower than initial P posteriors (expected due to degeneracies).
  - Posteriors for late-time evolution do not overlap causing bimodality (hinting at missing info/physics).

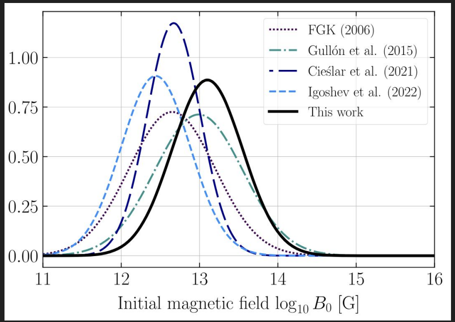

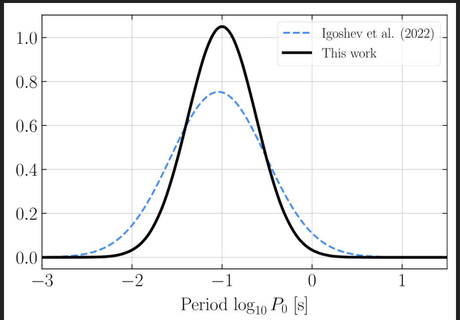

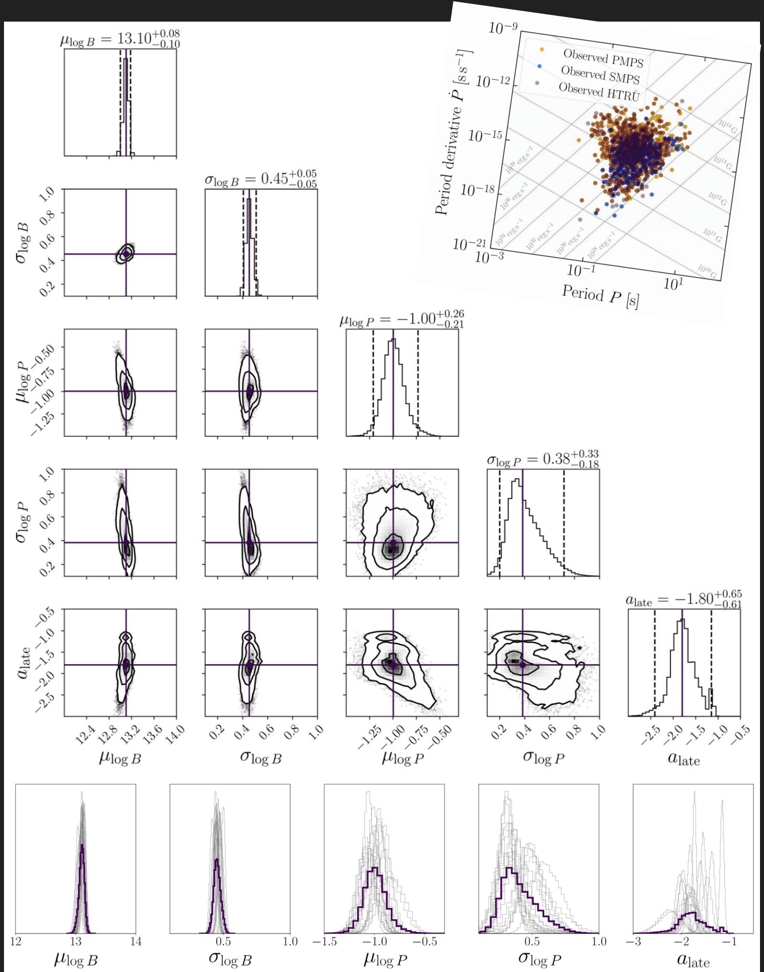

Simulator in Pardo et al. (2025)

- Using a **new radio luminosity**, we add two additional free parameters: the exponent, α, and the mean, μ_L, of the log-normally distributed normalisation factor L_0.

$$L_{\rm int} = L_0 \left(\frac{\dot{E}_{\rm rot}}{\dot{E}_{0,\rm rot}} \right)^{\alpha} ,$$

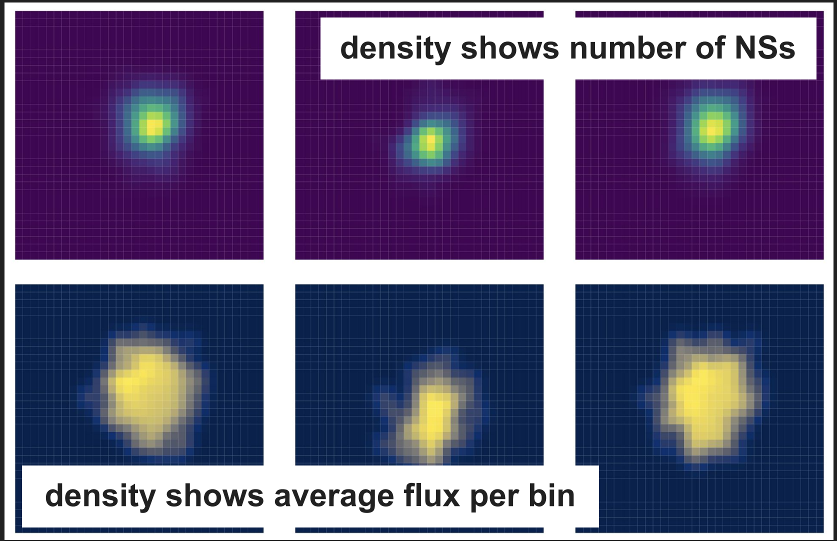

- In addition to our period-period derivative diagrams, we provide consistent radio flux measurements observed with the MeerKat telescope (Posselt et al. 2023) to our neural network. We convert these into similar density maps.

For 7 parameters, we can no longer use an amortised NPE approach.

TSNPE Pardo et al. (2025)

- We use the Truncated Sequential NPE approach (Deistler et al. 2022) to **focus on the regions of the parameter space** that match our observed data.

NPE

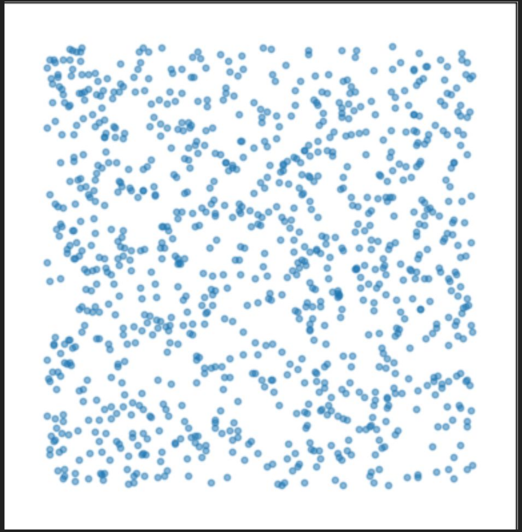

TSNPE

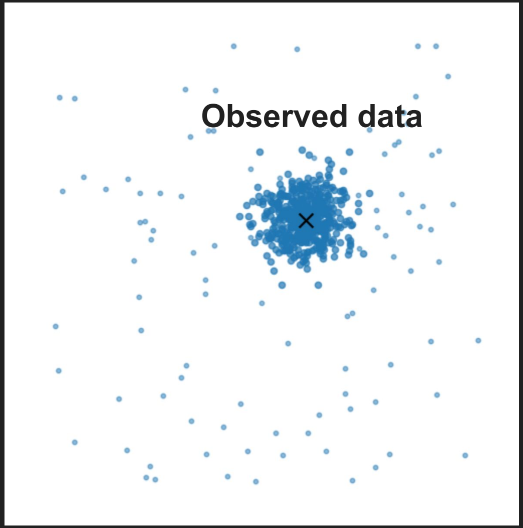

- We train over **10 rounds** producing 1,000 simulations in each round. This gives a total of 10,000 simulations.

- We use the baseline network from Graber et al. (2024) and train it 5 times with different initialisations to obtain an **ensemble prediction**.

Results III

- When evaluating the algorithm for a test sample with known ground truths, we find that we require roughly 8 rounds to obtain converged posteriors.

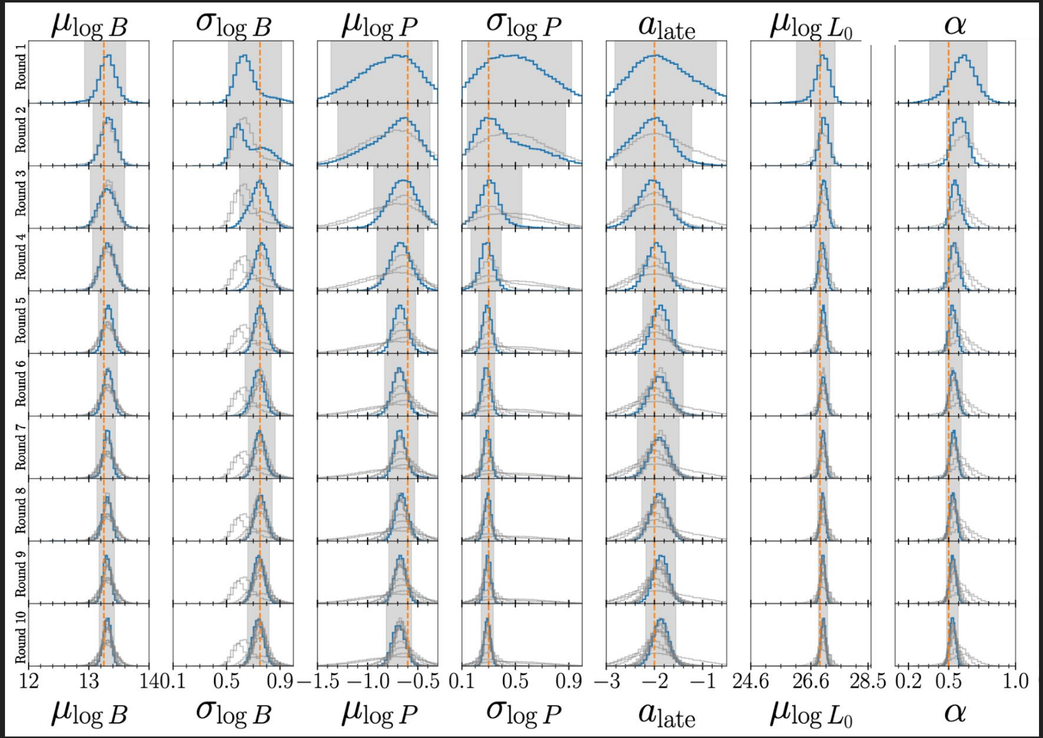

TSNPE is much more efficient and we require only 8,000 simulations compared to the 360,000 for NPE.

Vanessa.Graber@rhul.ac.uk

Results Pardo et al. (2025) IV

- When applying TSNPE on the observed population, we successfully infer 5 magneto-rotational and 2 luminosity parameters.

- We find that **adding flux information** fixes the issue of the bimodality in the late-time evolution seen in Graber et al. (2024).

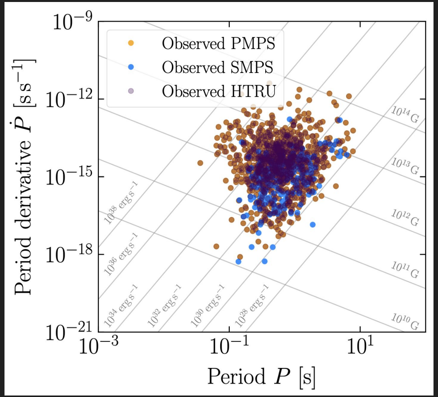

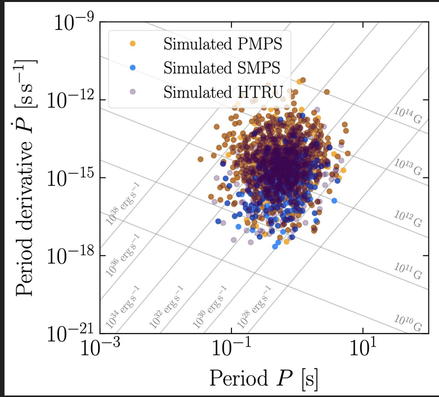

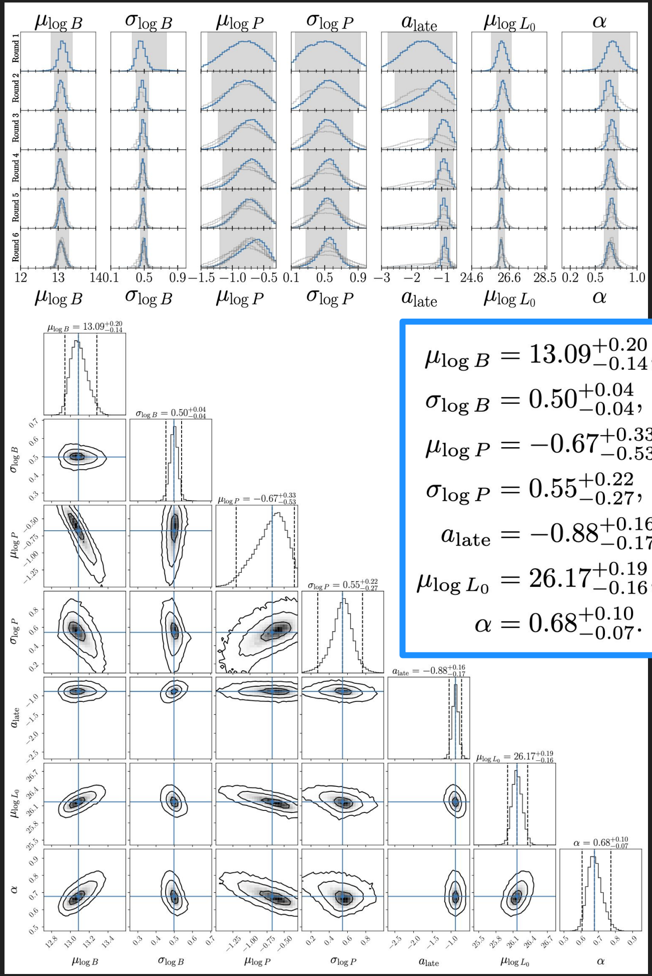

$$\mu_{\log B} = 13.09^{+0.20}_{-0.14},$$
$$\sigma_{\log B} = 0.50^{+0.04}_{-0.04},$$
$$\mu_{\log P} = -0.67^{+0.33}_{-0.53},$$
$$\sigma_{\log P} = 0.55^{+0.22}_{-0.27},$$
$$a_{\rm late} = -0.88^{+0.16}_{-0.17},$$
$$\mu_{\log L_0} = 26.17^{+0.19}_{-0.16},$$
$$\alpha = 0.68^{+0.10}_{-0.07}.$$

Outline

- Neutron stars
- Pulsar population synthesis
- Machine learning and sbi
- Inference results
- **Summary**

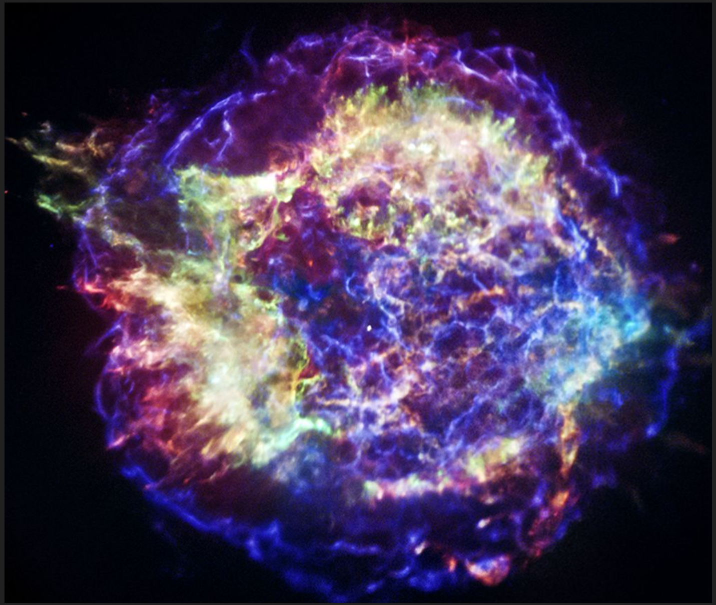

Cassiopeia A supernova remnant
(credit: NASA/CXC/SAO)

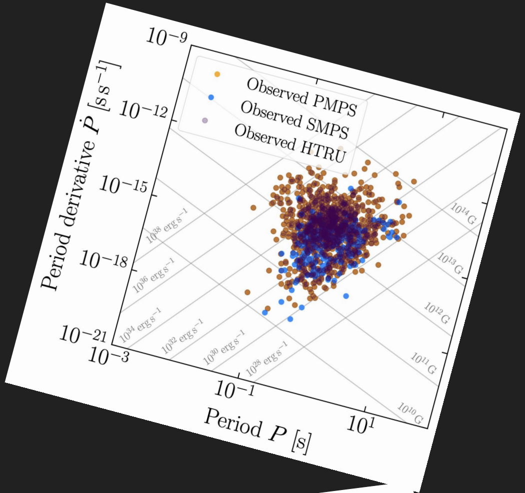

Summary

- Pulsar population synthesis bridges the gap between 3,500 observed pulsars and the true population.

- SBI with neural density estimators is a powerful tool to infer pulsar parameters. TSNPE is particularly efficient.

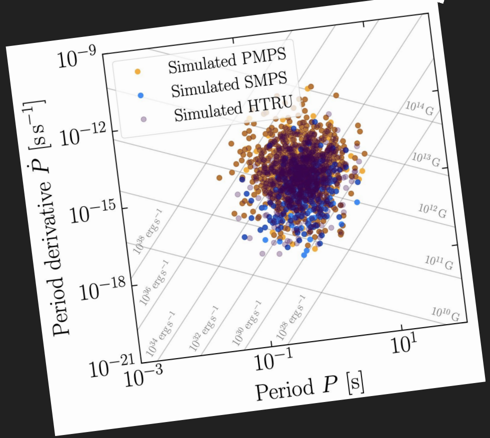

- Realistic simulators are very complex, making standard Bayesian inference impossible.

- We are now working to incorporate complementary X-ray observations and doing model comparison with SBI.

THANK YOU

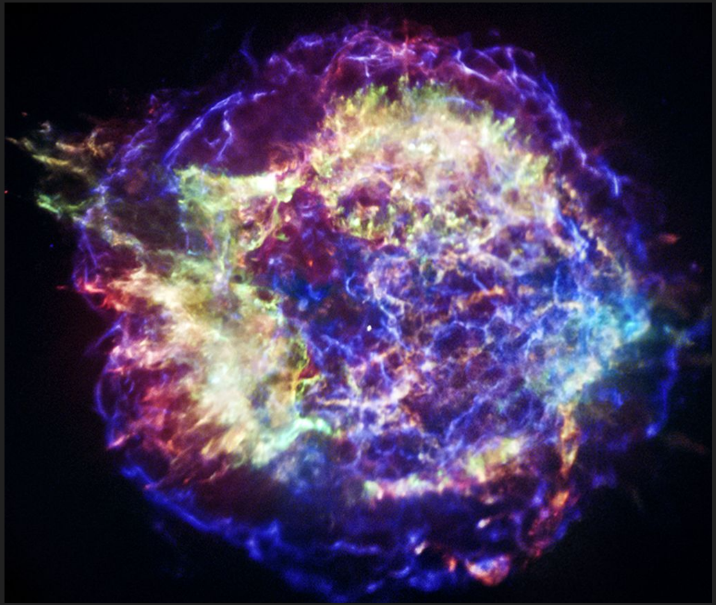

Cassiopeia A supernova remnant
(credit: NASA/CXC/SAO)