Draft version June 13, 2026
Typeset using LaTeX **twocolumn** style in AASTeX631

Constraints on the birth magnetic field distribution of isolated neutron stars from a unified population synthesis of radio pulsars, magnetars and XDINSs

Michele Ronchi [1,2,3] Celsa Pardo Araujo [2,3] Vanessa Graber [4] Nanda Rea [2,3] Clara Dehman [5] Davide De Grandis [2,3] Alessio Marino [2,3,6,7,8] Francesco Coti Zelati [2,3] Joeri van Leeuwen [1,9,10] and José A. Pons [5]

[1] *ASTRON, the Netherlands Institute for Radio Astronomy, Oude Hoogeveensedijk 4, 7991 PD Dwingeloo, The Netherlands*
[2] *Institute of Space Sciences (CSIC-ICE), Campus UAB, Carrer de Can Magrans s/n, 08193, Barcelona, Spain*
[3] *Institut d'Estudis Espacials de Catalunya (IEEC), Carrer Gran Capità 2–4, 08034 Barcelona, Spain*
[4] *Department of Physics, Royal Holloway, University of London, Egham, TW20 0EX, UK*
[5] *Departament de Física, Universitat d'Alacant, Ap. Correus 99, E-03080, Alacant, Spain*
[6] *Departamento de Física, Universidad de Santiago de Chile (USACH), Av. Víctor Jara 3493, Estación Central, Chile*
[7] *Center for Interdisciplinary Research in Astrophysics and Space Sciences (CIRAS), Universidad de Santiago de Chile*
[8] *INAF - IASF Palermo, via Ugo La Malfa 153, I-90146 - Palermo, Italy*
[9] *Leiden Observatory, Leiden University, Einsteinweg 55, 2333 CC Leiden, The Netherlands*
[10] *Leiden Institute for Advanced Computer Science (LIACS), Leiden University, Einsteinweg 55, 2333 CC Leiden, The Netherlands*

Submitted to ApJ

ABSTRACT

Understanding the Galactic population of isolated neutron stars within a unified framework provides key insights into their birth properties, evolutionary pathways, and the connections between different neutron star classes. In this work, we aim to reproduce both the observed radio and quiescent X-ray emission from isolated neutron stars, including radio pulsars, magnetars and X-ray dim isolated neutron stars (XDINSs). We develop a comprehensive population synthesis framework that models neutron star birth properties, their dynamical, rotational, and magneto-thermal evolution, as well as radio and X-ray emission, and the selection effects of corresponding surveys. To simulate realistic X-ray spectra, we account for magnetospheric resonant cyclotron scattering and interstellar absorption. Additionally, we model the observational bias introduced by magnetar outbursts, by linking the outburst rate to magnetic stresses in the stellar crust. We then employ a simulation-based inference method, namely truncated sequential neural posterior estimation, to reconstruct the birth properties, such as the initial magnetic field distribution. **We find that the Galactic neutron star population is best described by a two-component log-normal distribution of birth magnetic fields, with peaks at 5×10^{12} G and 10^{14} G, respectively. We further find that the Galactic neutron star birth rate should be around 3–4 per century, of which up to 50% are born as magnetars. Our results helps to quantify the contribution of neutron stars as central engines of powerful astrophysical transient phenomena, such as fast radio bursts, super-luminous supernovae and gamma ray bursts.**

Keywords: Machine learning — Neutron stars(1108) — Population synthesis — Pulsars(1306) — Radio pulsars(1353) —Simulation-based inference

Corresponding author: Michele Ronchi
ronchi@astron.nl

1. INTRODUCTION

Neutron stars are the compact remnants left over after the core-collapse supernova of massive stars. By combining extraordinary properties such as extreme gravity, ultra-strong magnetic fields, and rapid rotation, neutron

stars represent unique laboratories to study the behavior of matter and radiation in extreme environments. They have been observed to emit across the entire electromagnetic spectrum from radio to gamma-rays and showcase diverse emission behavior, from regular periodic pulses to more sporadic highly energetic bursts. Due to this variety of properties and behavior neutron stars have been classified into different classes (see, e.g., Kaspi 2010; Harding 2013; Borghese & Esposito 2023; Popov 2023). The bulk of known neutron stars have been discovered as isolated radio pulsars with estimated dipolar magnetic fields in the range $\sim 10^{12} - 10^{13}$ G. Decades of radio pulsar surveys have produced a statistically rich sample, now comprising over four thousand objects (see ATNF catalogue[1], Manchester et al. 2005). Among other classes of isolated neutron stars there are magnetars and X-ray Dim Isolated Neutron Stars (XDINSs). Magnetars are the neutron stars with the strongest magnetic fields exceeding 10^{14} G. Both their persistent and bursting emission is believed to be powered by the dissipation and instability of these strong magnetic fields. (see Turolla et al. 2015; Kaspi & Beloborodov 2017; Esposito et al. 2021; Rea & De Grandis 2025, for a review). XDINSs, also known as the "Magnificent Seven", are a class of isolated neutron stars with faint X-ray emission, showing an almost perfect thermal spectrum with broad absorption features (see Turolla 2009, for a review). Owing to their low luminosities, they are detectable only within a few hundred parsecs, which is why the currently known sources have estimated distances lower than 500 pc. If considered unrelated to one another, the existence of several neutron star classes challenges the limits of the core-collapse supernova rate in our Galaxy (Keane & Kramer 2008; Rozwadowska et al. 2021; Pardo-Araujo et al. 2026). Furthermore, in recent years the borders between the different classes of isolated neutron stars have started to blur. For example, magnetar-like burst activity has been discovered in some rotation-powered pulsars (RPPs) (Gavriil et al. 2008; Archibald et al. 2016) whereas pulsed radio emission was detected from several magnetars in outburst (e.g. Camilo et al. 2006). Furthermore, magnetar-like and faint radio emission have been detected from yet another class, so-called central compact objects (CCOs) (Rea et al. 2016; D'Aì et al. 2016; Zhang et al. 2025). These are young neutron stars found at the geometrical center of supernova remnants and have previously been detected only as stable X-ray emitters. These observational hybrid properties indicate that the different classes of isolated neutron stars are likely connected through evolutionary pathways, and therefore should be studied as a single population rather than as distinct, unrelated categories. A unified approach helps us to better understand the neutron stars birth properties, the relationship with their massive progenitors and supernova explosions that lead to their formation.

In general, the neutron star distribution in the spin period derivative, $\dot{P}$, *vs* spin period, P, plane encodes information about the initial spin and magnetic field distributions at birth and the evolutionary pathways of the population. Population synthesis represents a powerful computational framework used to simulate the birth, evolution, and observable properties of the populations of neutron stars in the Galaxy with the aim of comparing with the observed population and constraining physical models. Due to the rich radio pulsar sample, the main efforts of population synthesis studies have been focused on modeling the radio pulsar population alone (e.g. Narayan & Ostriker 1990; Faucher-Giguère & Kaspi 2006; Gullón et al. 2014; Cieślar et al. 2020; Graber et al. 2024; Pardo-Araujo et al. 2025). However, radio pulsars represent only a subset of the underlying neutron star population, and analyses based solely on radio detections may bias the inferred birth properties.

Both magnetars and XDINSs probe regimes of magnetic field strengths and evolutionary timescales poorly sampled by the radio pulsar population. Magnetars trace the young high-field end of the birth distribution and exhibit strong magnetic field decay, whereas XDINSs likely represent an evolved population of middle-aged, cooling neutron stars and are believed to be old magnetars (Viganò et al. 2013). Together, these thermally emitting sources offer crucial leverage on the shape and width of the initial magnetic field distribution, on the role of magnetic field decay, and on the connection between different neutron star classes. In particular, understanding the birth magnetic field distribution gives important clues on the formation mechanism during supernova explosions and the fraction of neutron stars that are born as magnetars (Makarenko et al. 2021; Pardo-Araujo et al. 2026). This in turn has important implications for studying the connections between neutron stars and powerful transients events such as fast radio bursts (FRBs) and gamma-ray bursts (GRBs) (see, e.g. Rea et al. 2015; Stratta et al. 2018; Beniamini et al. 2025; Zhang & Hu 2025).

Recent advances in magneto-thermal models (Viganò et al. 2021; De Grandis et al. 2021; Dehman et al. 2023c,b; Ascenzi et al. 2024) provide an improved theoretical basis for studying the evolutionary links between these classes of neutron stars and interpreting them in a unified framework. Incorporating realistic prescriptions

[1] https://www.atnf.csiro.au/research/pulsar/psrcat/

for the magneto-thermal evolution makes it possible to model not only radio pulsars but also X-ray–bright neutron stars within a single evolutionary scheme. Combining such models with population synthesis has been attempted in a few studies in the past two decades (e.g. Popov et al. 2010; Gullón et al. 2015). These studies showed that adding thermally emitting neutron stars helps to break degeneracies between simulation parameters and to better constrain their corresponding ranges, especially for those describing the initial magnetic field distribution. In particular, Gullón et al. (2015) showed that a single log-normal distribution for the initial magnetic fields is unlikely to explain both populations of radio pulsars and magnetars and an extra component extending to fields up to 5×10^{14} G is required.

In this work, we use the software `ML-Poppyns`[2] (Ronchi et al. 2021; Graber et al. 2024; Pardo-Araujo et al. 2025) to perform a comprehensive population synthesis study aimed at constraining the birth and evolutionary properties of isolated neutron stars by jointly considering the radio pulsar population and thermally emitting neutron stars, including magnetars and XDINSs. By combining realistic magneto-thermal evolutionary tracks with emission models in the radio and X-ray bands and survey-specific selection biases, we explore whether a unified scenario with a unique set of birth properties can account for the observed diversity of neutron stars. This multi-population, multi-wavelength approach provides a more complete picture of neutron star birth properties and offers new insights into the connections between distinct observational classes.

The paper is structured in the following way: in Section 2, we describe the observed dataset of radio pulsars and X-ray emitting neutron stars used in this work; in Section 3, we explain the ingredients of the `ML-Poppyns` population synthesis framework; in Section 4, we highlight the simulation-based inference algorithm used to perform the parameter inference; in Section 5, we describe the results and in Section 6, we discuss their implication and compare with other works.

2. OBSERVED NEUTRON STAR POPULATION

2.1. *Radio pulsars*

We consider the same observational sample as described in Pardo-Araujo et al. (2025), consisting of the radio pulsars detected by three major surveys conducted with Murriyang, the Parkes radio telescope: the Parkes Multibeam Pulsar Survey (PMPS) (Manchester et al. 2001; Lorimer et al. 2006), the Swinburne Intermediate-latitude Pulsar Survey (SMPS) (Edwards et al. 2001; Jacoby et al. 2009), and the low- and mid-latitude High Time Resolution Universe (HTRU) surveys (Keith et al. 2010). For the timing properties, i.e., P and $\dot{P}$, we rely on the ATNF catalogue version 2.6.0. For the radio fluxes we consider the data from the Thousand Pulsar Array (TPA) program (Johnston et al. 2020), which is part of the large survey project MeerTIME on the MeerKAT telescope. The TPA provides a consistently observed sample with well-calibrated flux measurements at 1.429 GHz as reported by Posselt et al. (2023). The number of pulsars detected by these three surveys are the following:

- PMPS: 1045,
- SMPS: 218,
- HTRU: 1095.

The discrepancy between the number for the HTRU survey reported here and in Pardo-Araujo et al. (2025) is due to reprocessing of archival data that leads to the discovery of 58 new isolated pulsars (Sengar et al. 2025). For more details on the filters applied to construct this observational dataset we refer to Pardo-Araujo et al. (2025).

2.2. *Magnetars, XDINSs and other thermally emitting X-ray pulsars*

The observational sample we consider includes all neutron stars that have a statistically significant thermal component in the soft X-ray band in the energy range 0.1 - 10 keV (see Table 1). The origin of this thermal emission is attributed to the residual heat stored in the neutron star interior and to the Ohmic dissipation of the magnetic field in the neutron star crust (Viganò et al. 2013, 2021). We only considered archival *Chandra* and/or *XMM-Newton* observations, as they provide the best combination of effective area and angular resolution among past and present X-ray observatories. For all sources, we estimated the absorbed flux in the energy range 0.1 - 10 keV. More details on the data reduction and analysis can be found in Coti Zelati et al. (2018); Marino et al. (2024), Dehman et al., in prep.. The resulting observational dataset contains 23 magnetars and 7 XDINSs (1 of which is missing the measurement of the spin period derivative, $\dot{P}$).

Magnetars have been mainly detected through episodes of very energetic outburst emission. However, here we only consider data during their quiescence phase where their emission is mostly of thermal origin and can be described by one or more black-body components possibly with a power-law tail at higher energies. The

[2] https://ice-csic-astroexotic.github.io/code/ml_poppyns/

Table 1. Neutron stars with a significant thermal component in the quiescence phase used in this work, classified as magnetars and XDINSs. We report their spin period, P, spin-period derivative, $\dot{P}$, and the absorbed flux, $S_{\rm X,abs}$, in the energy range 0.1 - 10 keV. Fluxes values have to be considered with a 10% relative error.

Source	P [s]	$\dot{P}$ [10^{-11} s s^{-1}]	$S_{\rm X,abs}$ [erg s^{-1} cm^{-2}]	class
SGR1627-41	2.59	1.9	4.20×10^{-14}	magnetar
1E2259+586	6.98	0.048	3.55×10^{-11}	magnetar
XTEJ1810-197	5.54	0.283	5.30×10^{-13}	magnetar
SGR1806-20	7.75	7.5	5.49×10^{-12}	magnetar
CXOUJ1647-4552	10.61	0.097	8.00×10^{-13}	magnetar
SGRJ0501+4516	5.76	0.594	2.50×10^{-12}	magnetar
1E1547-5408	2.07	4.77	3.20×10^{-13}	magnetar
SGRJ0418+5729	9.08	0.0004	1.00×10^{-14}	magnetar
SGRJ1833-0832	7.57	0.35	6.00×10^{-14}	magnetar
SwiftJ1822.3-1606	8.44	0.013	2.30×10^{-13}	magnetar
SwiftJ1834.9-0846	2.48	0.806	1.00×10^{-14}	magnetar
1E1048.1-5937	6.46	2.18	5.56×10^{-12}	magnetar
SGRJ1745-2900	3.76	3.06	2.00×10^{-14}	magnetar
SGRJ1935+2154	3.24	1.43	8.60×10^{-13}	magnetar
1E1841-045	11.79	4.09	2.33×10^{-11}	magnetar
SGR1900+14	5.20	9.2	3.92×10^{-12}	magnetar
4U0142+614	8.69	0.2	1.13×10^{-10}	magnetar
1RXSJ170849.0-4009	11.01	1.95	3.75×10^{-11}	magnetar
CXOUJ171405.7-3810	3.83	6.4	1.69×10^{-12}	magnetar
PSRJ1119-6127	0.407	0.4	4.80×10^{-14}	magnetar
PSRJ1622-4950	4.326	1.7	9.00×10^{-15}	magnetar
SwiftJ1818.0-1607	1.36	5.0	2.51×10^{-14}	magnetar
3XMMJ1852+0033	11.559	0.014	4.46×10^{-13}	magnetar
RXJ0420.0-5022	3.45	0.002759	5.01×10^{-13}	XDINS
RXJ1856.5-3754	7.055	0.003	2.00×10^{-11}	XDINS
RXJ2143.0+0654	9.428	0.0041	2.51×10^{-12}	XDINS
RXJ0720.4-3125	8.391	0.006983	1.00×10^{-11}	XDINS
RXJ0806.4-4123	11.37	0.0055	2.51×10^{-12}	XDINS
RXJ1308.6+2127	10.31	0.011	3.16×10^{-12}	XDINS
RXJ1605.3+3249	3.39	-	7.94×10^{-12}	XDINS

latter has been attributed to the resonant cyclotron scattering of thermal seed photons emitted by the surface as they interact with the charged particles gyrating around the magnetic field lines in the magnetosphere (see Section 3.4, Rea et al. 2008; Zane et al. 2009; Beloborodov 2013).

We do not include CCOs in this work since our simulation pipeline does not explicitly account for fallback accretion that could lead to the formation of objects like CCOs (Viganò & Pons 2012).

Finally, we not that also RPPs can manifest X-ray spectra with both a thermal and a non-thermal component. Especially for the youngest ones, these two components are usually attributed to the release of heat from the surface and to the synchro-curvature emission from charged particles in the magnetosphere, respectively (Becker & Truemper 1997; Xu et al. 2025). In this work, we exclude RPPs as many of them have spectra contaminated by the contribution of a pulsar-wind nebula (e.g., Gotthelf 2003; Cheng et al. 2004). As the nebula is brighter in X-rays than the central neutron star, this introduces a detection bias which we are not modeling in our simulations. Moreover, contrary to magnetars and XDINSs, most RPPs are already detected in radio and probe a range of magnetic fields which is already well-represented by the rich sample of radio pulsars.

3. POPULATION SYNTHESIS

To generate a synthetic population of neutron stars, we performed Monte Carlo simulations to model both the dynamical and magneto-rotational evolution of neutron stars. We then model their radio and X-ray emis-

sion and apply observational biases both for the radio and X-ray surveys to compare the simulated populations with observations. We follow the same approach as in Ronchi et al. (2021); Graber et al. (2024); Pardo-Araujo et al. (2025), and only provide a brief summary of the ingredients referring the reader to these works for more detail on the employed methodology.

3.1. *Dynamical evolution*

Following the same approach as in Pardo-Araujo et al. (2025), we assume the dynamical properties are decoupled from the magneto-rotational ones, allowing us to create a database containing the information of a dynamically evolved population of neutron stars. With this purpose, we simulated a population of 2×10^7 neutron stars, assigning them uniformly distributed random ages up to 10^8 yr to ensure a sufficiently large database for the subsequent steps.

For the initial positions we use the same strategy outlined in Ronchi et al. (2021) (but see also Graber et al. 2024; Pardo-Araujo et al. 2025). We assume that the progenitor OB stars follow the spiral arms parametrized with a logarithmic model as in Yao et al. (2017) and the Galactocentric exponential radial distribution found by Verberne & Vink (2021) for supernova remnants. **We also take into account the local arm (see Ronchi et al. 2021) which is important to model the Sun neighbourhood, especially for reproducing the XDINSs population**. For the Galactic height, we assume an exponential disk profile (Wainscoat et al. 1992) with a characteristic scale height of 0.18 kpc, consistent with the vertical distribution of young massive stars in the Milky Way (Li et al. 2019). The kick velocities are drawn from the log-normal distribution found in Disberg & Mandel (2025) with mean of 5.6 and standard deviation of 0.68. After setting the initial conditions, we evolve the neutron stars' positions and velocities in time by solving the Newtonian equation of motion with the same prescription as in Graber et al. (2024); Pardo-Araujo et al. (2025). In this way, we obtain a dynamically evolved database of neutron stars that we can sample from to perform the following steps of the simulation, i.e., the magneto-rotational evolution and the detection.

3.2. *Magneto-rotational evolution*

In order to model the magneto-rotational evolution of neutron stars we assume that the initial spin periods follow a log-normal distribution of the form:

$$\mathcal{P}(\log P_0) = \frac{1}{\sqrt{2\pi}\sigma_{\log P}} \exp\left(-\frac{(\log P_0 - \mu_{\log P})^2}{2\sigma^2_{\log P}}\right). \tag{1}$$

Here, and in the following, we use the subscript log to refer to $\log_{10}$ to not clutter the notation.

For the initial magnetic field distribution, we consider a double log-normal parametrized in the following way:

$$\begin{aligned}\mathcal{P}(\log B_0) = w\mathcal{N}(\log B_0, \mu_{\log B,1}, \sigma_{\log B,1}) \\ + (1-w)\mathcal{N}(\log B_0, \mu_{\log B,2}, \sigma_{\log B,2}),\end{aligned} \tag{2}$$

where

$$\mathcal{N}(\log B_0, \mu_{\log B,i}, \sigma_{\log B,i}) = \frac{1}{\sqrt{2\pi}\sigma_{\log B,i}} \exp\left(-\frac{(\log B_0 - \mu_{\log B,i})^2}{2\sigma^2_{\log B,i}}\right), \tag{3}$$

with $i = 1, 2$ and w representing a weight parameter for the first component with a range between 0 and 1.

The choice of this model for the initial magnetic field distribution is given by the fact that a distribution extended to stronger magnetic fields is required to explain the population of magnetars (see Popov et al. 2010; Gullón et al. 2015; Sautron et al. 2025). The presence of two components in the initial magnetic field distribution implicitly assumes that the birth magnetic field could originate either from different progenitors or mechanisms that enhance the magnetic field during the core collapse (Duncan & Thompson 1992; Barrère et al. 2022).

The initial inclination angle χ between the spin axis and the magnetic field axis is drawn randomly from a uniform distribution in spherical coordinates, i.e., with a probability $\mathcal{P}(\chi) = \sin\chi$.

After establishing the initial conditions, the spin period and the inclination angle are evolved in time by solving the coupled differential equations (Spitkovsky 2006; Philippov et al. 2014):

$$\dot{P} = \frac{\pi^2}{c^3}\frac{B^2R^6}{IP}\left(\kappa_0 + \kappa_1 \sin^2\chi\right), \tag{4}$$

$$\dot{\chi} = -\frac{\pi^2}{c^3}\frac{B^2R^6}{IP^2}\left(\kappa_2 \sin\chi\cos\chi\right), \tag{5}$$

where c is the speed of light, $R = 12.59$ km is the neutron star radius for a fiducial neutron star mass $M = 1.4\,M_\odot$ assuming the equation of state BSk24 (Pearson et al. 2018), and $I \simeq 2MR^2/5 \approx 1.78 \times 10^{45}\,\mathrm{g\,cm^2}$ is the stellar moment of inertia. For realistic pulsars surrounded by plasma-filled magnetospheres, we choose $\kappa_0 = \kappa_1 = \kappa_2 = 1$.

3.3. *Magneto-thermal evolution models*

To model the coupled evolution of magnetic field and thermal luminosity of neutron stars, we rely on the results of 2D magneto-thermal simulations (see Viganò

et al. 2021). We assumed the equation of state BSk24 (Pearson et al. 2018), which has been shown to be able to explain the luminosities of thermally emitting neutron stars (Marino et al. 2024), a neutron star mass $M_{\rm NS} = 1.4\ {\rm M}_\odot$ and a corresponding radius $R_{\rm NS} = 12.59\,{\rm km}$. The impurity parameter in the pasta layer is fixed to 100 (Pons et al. 2013). For the impurity in the outer and inner crust (excluding the pasta layer), the fits of Carreau et al. (2020) have been used (see Fig. 5 in that paper). The blanketing envelope model we adopt is that of Potekhin et al. (2015), composed of heavy elements such as iron and accounting for magnetic-field effects. For a discussion of the impact of different envelope compositions and the magnetization of the envelope on magneto-thermal evolution, we refer the reader to Dehman et al. (2023a). Moreover, we consider a crust-confined magnetic field configuration that includes only the dipolar ($\ell = 1$) poloidal and toroidal magnetic field components. We set the two components to have the same maximum magnetic strength. As a result, the poloidal dipole component contains $\sim 90\%$ of the total magnetic energy, while the remaining $\sim 10\%$ are stored in the toroidal component. In the top panel of Fig. 1, we show the poloidal dipolar magnetic field evolution with an analytical parametrization for magnetic field evolution (dashed lines). As current magneto-thermal simulations can only properly model the surface temperature and field evolution until $\sim 10^6$ yrs, in this work we used the same parametrization as described in Graber et al. (2024). This approach captures the trend of field decay for different initial magnetic field strengths by combining several broken power laws together with a late-time power-law evolution with power-law index $a_{\rm late}$. Furthermore to avoid the field decaying to unrealistically small numbers at very late times, we sample the $\log_{10}$ of the final fields from a Gaussian distribution with mean $\mu = 8.5$ and standard deviation $\sigma = 0.5$ in line with the distribution seen for old millisecond pulsars (see Appendix A in Graber et al. 2024, for more details). This allows us to easily extract the dipolar field strength, B, at different times, t, to compute the magneto-rotational evolution of our synthetic pulsars.

Together with the magnetic field decay, magneto-thermal simulations provide the coupled evolution in time of the temperature profile, $T(\theta, t)$, on the neutron star surface, where θ denotes the polar angular coordinate. Given a temperature profile, $T(\theta, t)$, at a given time, we can compute the total luminosity that a distant observer would estimate by assuming black-body emission from every element of surface area and integrating over the neutron star surface:

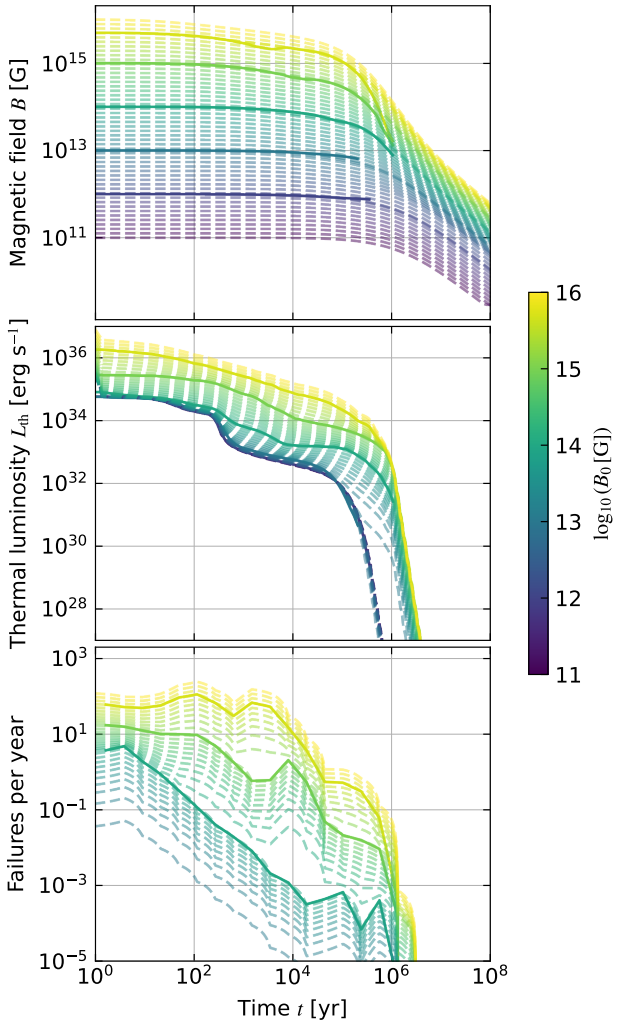

Figure 1. *Top panel*: Evolution of the poloidal dipolar B-field. The five solid curves represent realistic two-dimensional simulations of magneto-thermal evolution in the neutron-star crust (Viganò et al. 2021). The dashed lines represent the evolution predicted by the analytical prescription used in Graber et al. (2024) with the parameter $a_{\rm late} = -1.0$. *Middle panel:* The corresponding evolution of the surface thermal luminosity. The five solid curves represent realistic two-dimensional simulations of magneto-thermal evolution in the neutron-star crust (Viganò et al. 2021). The dashed lines represents the evolution predicted by interpolating the original curve with a bivariate spline and extrapolating to a range of initial dipolar magnetic field values. For $t > 10^6$ yr we assume a power-law evolution $\propto t^{-10}$ (see text). *Bottom panel:* The corresponding evolution of the rate of crust failures due to magnetic stresses. The three solid curves represent the predicted crust-failure rate from realistic two-dimensional simulations of magneto-thermal evolution in the neutron-star crust (Viganò et al. 2021). The dashed lines represents the evolution predicted by interpolating the original curves with a bivariate spline and extrapolating to a range of initial dipolar magnetic field values. In all panels, the colors represent the initial poloidal dipolar magnetic-field strength, B_0.

$$L_{\rm th}(t) = 2\pi \int_0^{\pi} \sigma_{\rm SB} T_\infty(\theta, t)^4 R_\infty^2 \sin\theta d\theta. \quad (6)$$

Here, $\sigma_{\rm SB} = 5.67 \times 10^{-5}\,{\rm erg\,cm^{-2}\,K^{-4}\,s^{-1}}$ is the Stefan-Boltzmann constant, and T_∞ and R_∞ denote the temperature and neutron star radius respectively that a distant observer would measure due to the curvature of space-time around the neutron star. These are defined as:

$$T_\infty = \left(1 - \frac{2GM_{\rm NS}}{c^2 R_{\rm NS}}\right)^{1/2} T, \quad (7)$$

$$R_\infty = \left(1 - \frac{2GM_{\rm NS}}{c^2 R_{\rm NS}}\right)^{-1/2} R_{\rm NS}. \quad (8)$$

When the temperature becomes too low, the routines that model the microphysics in magneto-thermal simulations no longer provide reliable results. This happens at ages between $\sim 10^5 - 10^6$ yr. To perform population synthesis simulations at late times, we need to extend the luminosity after $\sim 10^6$ yr by considering the cooling during the photon dominated era. In order to get an estimate of the thermal evolution at the late stage, we assume that the surface temperature of the envelope, $T_{\rm e}$, is linked to the temperature at the bottom of the envelope, $T_{\rm b}$, at the interface with the crust via the relation $T_{\rm e} \sim T_{\rm b}^{0.5+\alpha}$, where $\alpha \sim 0.05$ (see Eq. 34 in Page et al. 2004). During the photon cooling era, we have that $T_{\rm e} \propto t^{-\frac{1}{8\alpha}}$ (see Eq. 42 in Page et al. 2004), which translates into a luminosity evolution $L_\gamma \propto T_{\rm e}^4 \propto t^{-\frac{1}{2\alpha}} = t^{-10}$ (see Eq. 38 in Page et al. 2004). Therefore for the late time evolution of the luminosity, we consider a power-law with the trend $L_{\rm th} \propto t^{-10}$. In the middle panel of Fig. 1, we show the resulting surface thermal luminosity evolution. As we performed simulations only for discrete values of the initial poloidal dipolar magnetic fields, we interpolate using a bivariate spline in order to have an estimate of the luminosity for any initial dipolar magnetic field. This approximation is used to extrapolate the luminosity to all initial dipolar magnetic field values in the range $10^{11} - 10^{16}$ G. These interpolated curves allow us to obtain the thermal luminosity of a neutron star given its initial poloidal dipolar magnetic field strength and its age.

3.4. *Modeling the X-ray thermal emission*

Knowing the thermal luminosity at the surface, we can compute an average effective surface temperature as measured by a distant observer:

$$\bar{T}_\infty = \left(\frac{L_{\rm th}}{4\pi R_\infty^2 \sigma_{\rm SB}}\right)^{1/4}. \quad (9)$$

This temperature can be used to define a black-body spectrum.

Magnetar spectra show a non-thermal tail that is attributed to the resonant cyclotron scattering effect (Lyutikov & Gavriil 2006; Rea et al. 2008; Nobili et al. 2008). In this process the seed thermal photons coming from the neutron star surface interact with the magnetospheric plasma composed mainly of electrons and positrons gyrating along the magnetic field lines. Photons that have frequencies, ω, close to the cyclotron frequency $\omega_B = eB/(m_e c)$, will experience resonant scattering and their energy will be boosted. Multiple scattering events distort the black-body spectrum and introduce a power-law tail at high energies.

To compute the intensity spectrum distorted by resonant cyclotron scattering, $I_{\rm RCS}(E)$, where $E = h\nu = \hbar\omega$ is the energy of the photons, we used the simplified semi-analytical 1D model described in Lyutikov & Gavriil (2006) (see Appendix A for more details).

Once the scattered spectrum has been computed, we estimate the total luminosity through the integral:

$$L_{\rm X}(E) = 4\pi R_\infty^2 \int_0^{2\pi} \int_0^{\pi/2} I_{\rm RCS}(E) \cos\theta' \sin\theta' d\theta' {\rm d}\phi' \quad (10)$$

$$= 4\pi R_\infty^2 \pi I_{\rm RCS}(E),$$

where θ' denotes the angle between the normal to the neutron star surface and a random line of sight direction and ϕ' is the azimuthal angle. By knowing the distance d from the neutron star we can then derive the intrinsic flux density as:

$$S_{\rm X}(E) = \frac{L_{\rm X}(E)}{4\pi d^2} = \pi \left(\frac{R_\infty}{d}\right)^2 I_{\rm RCS}(E). \quad (11)$$

In order to compute the X-ray flux reaching the Earth, we need to correct for the interstellar absorption due to photoionization. Given an intrinsic (not absorbed) X-ray flux density, $S_{\rm X}(E)$, where E denotes the energy in keV, the absorbed flux can be derived by (see Wilms et al. 2000):

$$S_{\rm X,obs}(E) = e^{-\sigma_{\rm ISM}(E) N_{\rm H}} S_{\rm X}(E). \quad (12)$$

Here, $\sigma_{\rm ISM}(E)$ in units of cm^2 per hydrogen atom represents the energy-dependent effective absorption cross section of the interstellar medium (ISM), which is normalized on the hydrogen atom number. Taking into account the phases of the ISM, this cross section can be written as:

$$\sigma_{\rm ISM}(E) = \sigma_{\rm gas}(E) + \sigma_{\rm molecules}(E) + \sigma_{\rm grains}(E). \quad (13)$$

To compute $\sigma_{\rm ISM}(E)$, we will only take into account the neutral atomic gas, that is the $\sigma_{\rm gas}(E)$ term, ne-

glecting the possibility of ionization and the presence of molecules and grains in the ISM. As shown in Wilms et al. (2000) (see also Willingale et al. 2013), including these effects would give only minor corrections to the total effective cross sections. Therefore, we can rewrite the equation above as (see Wilms et al. 2000):

$$\sigma_{\rm ISM}(E) \sim \sigma_{\rm gas}(E) \sim \sum_Z A_Z \sigma_{{\rm ph},Z}(E), \qquad (14)$$

where, for an element with atomic number Z, the relative abundance with respect to hydrogen is defined as $A_Z = N_Z/N_{\rm H}$. Table 2 in Wilms et al. (2000) reports the logarithm of the relative abundance $\log_{10}(N_Z/N_{\rm H}) + 12$, where by definition the abundance of hydrogen is set to 12. $\sigma_{{\rm ph},Z}(E)$ denotes the photoionization cross section. To derive the total effective cross section $\sigma_{\rm ISM}(E)$ we have translated the `Fortran 77` routines written by Balucinska-Church & McCammon (1992) into `Python`. The original routines implement polynomial fits of the atomic photoionization cross sections in the energy range of 0.03 – 10 keV for seventeen elements.

In Eq. (12), $N_{\rm H}$ denotes the hydrogen column density which has units of hydrogen atoms cm^{-2} and is defined as the integral of the spatial density of neutral hydrogen, $n_{\rm H}$, along the line of sight:

$$N_{\rm H} = \int n_{\rm H} dl, \qquad (15)$$

where l represents the path length along the line of sight. For a given neutron star with known equatorial sky coordinates (RA, DEC) and heliocentric distance, d, the value of $N_{\rm H}$ can be estimated using the reddening map of the Galaxy and the calibration factor provided by Doroshenko (2024). This factor converts reddening, $E(B-V)$, provided by the map into an $N_{\rm H}$ value assuming a set of abundances for the ISM. By default we use the abundances specified in Wilms et al. (2000).

Finally, to compute the total flux, we integrate the flux density in Eq. (12) in the energy band 0.1 – 10 keV.

3.5. *Crustal failures and outburst activity*

Magnetars have been primarily discovered through their flaring and bursting activity. This is often accompanied by sudden increases in emission by several orders of magnitude with respect to their persistent thermal and non-thermal emission that could last for months or years, a phenomenon commonly referred to as an outburst (Rea & Esposito 2011; Coti Zelati et al. 2018). This highly variable and powerful activity is likely linked to reconnection events in neutron stars magnetosphere, possibly triggered by the star's interior evolution. In particular, current models suggest that the evolution and instability of the ultra-strong magnetar magnetic field causes mechanical stresses in the neutron star crust. These accumulate to the point of exceeding the breaking strain of the crustal lattice, generating an event that can locally heat the matter and drag magnetic field lines along. The resulting perturbation can then propagate outwards triggering magnetospheric activity. (Beloborodov 2009; Chugunov & Horowitz 2010; Carrasco et al. 2019; Dehman et al. 2020; Lander 2023). As the flaring and outburst activity represents a crucial bias towards the discovery of highly magnetized neutron stars, we need to account for this in our model and predict the rate of these events for a given neutron star. In this regard, 2D magneto-thermal simulations also provide an estimate of when and where the crust will fail due to the stresses mentioned above (Perna & Pons 2011; Dehman et al. 2020). This information can then be used to compute the rate of crustal failures as a function of time and initial poloidal dipolar magnetic field strength during the magnetic field evolution. To evaluate the rate, we divide the total evolution time span into logarithmic bins and count how many failure events are produced in each bin. To obtain the rate of failures in each bin, we then divided the total count in each bin by the bin width.

In the bottom panel of Fig. 1, we show the rate of crust failures as a function of time predicted by our 2D simulations of magneto-thermal evolution. As we performed simulations only for discrete values of the initial dipolar magnetic fields, we interpolate using a bivariate spline in order to have an estimate of the failure rate for any initial dipolar magnetic field. This approximation is used to extrapolate the rate to all initial dipolar magnetic field values in the range $10^{11} - 10^{16}$ G. As for the thermal luminosity, these interpolated curves allow us to obtain the crust-failure rate of a neutron star given its initial poloidal dipolar magnetic field strength and age.

3.6. *Radio detection*

To compute the radio fluxes, beam geometry and propagation effects due to the interaction of the radio waves with the interstellar medium we use the same prescription as in Graber et al. (2024); Pardo-Araujo et al. (2025). In particular, to model the intrinsic radio luminosity we adopted Eq. (7) in Pardo-Araujo et al. (2025) where the mean of the logarithm of the radio luminosity, $\mu_{\log L_0}$, and the power-law index, α_L, determining the dependence on the spin-down power, are free parameters. As in Graber et al. (2024); Pardo-Araujo et al. (2025), we consider three major radio surveys conducted

with Murriyang, the Parkes radio telescope: the Parkes Multibeam Pulsar Survey (PMPS) (Manchester et al. 2001; Lorimer et al. 2006), the Swinburne Intermediate-latitude Pulsar Survey (SMPS) (Edwards et al. 2001; Jacoby et al. 2009), and the low- and mid-latitude High Time Resolution Universe (HTRU) surveys (Keith et al. 2010). A summary of all relevant survey parameters is provided in Table 1 in Graber et al. (2024). For each simulated neutron star whose radio beam intercepts our line of sight, we compute the signal-to-noise ratio using the radiometer equation (Lorimer & Kramer 2012):

$$S/N = \varepsilon \frac{S_{\rm mean} G \sqrt{n_{\rm pol} \Delta f_{\rm bw} t_{\rm obs}}}{\beta \left[T_{\rm sys} + T_{\rm sky}(l,b)\right]} \sqrt{\frac{P - w_{\rm obs}}{w_{\rm obs}}}. \quad (16)$$

For a more detailed description of all relevant parameters in this equation, we refer to Graber et al. (2024); Pardo-Araujo et al. (2025). We only stress here that compared to Graber et al. (2024); Pardo-Araujo et al. (2025), we have updated the radiometer equation by including the efficiency factor, ε, a function of the duty cycle $\delta = w_{\rm obs}/P$, as determined by Morello et al. (2020) (see their Eq. 44). Since pulsars are predominantly found via incoherent searches based on Fast Fourier Transforms (FFTs), this efficiency factor models the decay in sensitivity of pulsar searches when looking for pulsars with small duty cycles. Sometimes this efficiency factor is incorporated into the degradation factor, β, by considering values larger than ~ 1.25 (see discussion in Morello et al. 2020). Here we fix β to 1.25 assuming that it only accounts for imperfections during the digitization of the signal.

3.7. *Modeling the X-ray detection biases*

The observational sample for the X-ray emitting neutron stars is difficult to reproduce in a simulation framework as the sample is subject to complex observational biases which are not well under control. Many of these sources have been discovered through targeted observations with different X-ray instruments. Therefore, the sky coverage and corresponding threshold sensitivities are very inhomogeneous and difficult to reproduce in a simulation framework. Indeed, many magnetars have been discovered through their outburst activity whose high energy emission triggered all-sky X-ray monitors onboard Swift–BAT or Fermi–GBM. This has enabled follow-up campaign observations that allowed us to detect a periodicity and identify them as neutron stars (Rea & Esposito 2011; Coti Zelati et al. 2018). Moreover, in the observed catalog, for sources that underwent a magnetar-like outburst, we are only considering those with detected quiescent emission (note that some magnetars were identified as such during an outburst event but their quiescent emission is too faint to be detectable). On the other hand, sources like XDINSs have very stable X-ray emission over timescales of several decades, and despite their relatively faint X-ray luminosities can be detected due to their high fluxes, a consequence of their close-by distances. To try to encompass all these different observational biases, we adopt a simplified approach to model the X-ray detection.

First, we consider an all-sky coverage for X-ray surveys, i.e., we assume that X-ray instruments have scanned the entire sky. We assume that for fluxes in the energy range 0.1 - 10 keV, the observed sample is complete above a value of $\sim 10^{-11}\,{\rm erg\,s^{-1}\,cm^{-2}}$. When representing the observed X-ray population in the $\log N - \log S$ plane, where N represents the number of sources above a given flux value S (see Section 5), we observe a change in slope at fluxes lower than $10^{-12}\,{\rm erg\,s^{-1}\,cm^{-2}}$ (see black line in Fig. 4 and Fig. 8), meaning that we are probably starting to miss sources below this flux value. Therefore, we consider an average flux threshold of $10^{-12}\,{\rm erg\,s^{-1}\,cm^{-2}}$ in the X-ray band $0.1 - 10\,{\rm keV}$ with a standard deviation of 0.5 (in log). This assumption is also in agreement with averages flux thresholds of early all-sky surveys like with ROSAT (Truemper 1982; Voges et al. 1999). The implementation of this threshold allows the detection of all sources that are intrinsically very bright or close-by like the XDINSs.

On top of this first filter, we assume a second filter for all simulated magnetar-like sources that show outburst activity. In particular, for every simulated neutron star, we estimate the current crust-failure rate given its initial magnetic field and age as explained in Section 3.5. Given this rate, we estimate the probability for each star to experience a crust failure event that could have triggered an outburst in the last 30 years. This value is motivated by the fact that continuous all-sky monitoring capable of detecting new activity from magnetars has only been available for approximately the past 30 years, beginning with the launch of the Rossi X-ray Timing Explorer (RXTE) mission (Bradt et al. 1993). Prior to this, transient magnetar activity had been significantly harder to detect. Moreover, we assume that all outburst events are energetic enough to trigger a detection and deep follow-up observations with X-ray instruments. For the latter, we assume an average flux threshold of $10^{-14}\,{\rm erg\,s^{-1}\,cm^{-2}}$ with a standard deviation in log of 0.5. We assume this flux threshold in order to match the faintest detected magnetar (i.e., SGRJ0418+5729 in Table 1).

In general, the choice for the flux threshold values considered in this section comes from both trying to repro-

duce the observed X-ray flux distribution and the typical sensitivity of X-ray instruments for short and long exposure times (see for example Fig. 3 in Watson et al. 2001; Weisskopf et al. 2002; Gehrels et al. 2004).

3.8. *Representation of simulated output*

The output of the simulations, consisting of detected mock neutron stars in radio and X-rays, is represented in the form of 2D $P-\dot{P}$ maps following the same strategy as in Graber et al. (2024); Pardo-Araujo et al. (2025). In particular, for each simulated survey, we produce two maps: a density $P-\dot{P}$ map and an average flux map of the $P-\dot{P}$ diagram. Both types of maps have ranges set to $[0.01, 100]$ s and $[10^{-20}, 10^{-9}]$ s s^{-1} for the P and $\dot{P}$ axis, respectively, and a resolution of 32 bins along both axes. The density map contains information on the number of pulsars detected in each bin, while the flux maps store the average value of the radio flux in Jy or the X-ray flux in erg s^{-1} cm^{-2} of the neutron stars falling into that specific bin for the radio surveys or X-ray survey, respectively. Therefore the output of each simulation is summarized in a total of eight 2D maps, six for the radio surveys and two for the X-ray survey.

In Fig. 2, we show the maps for the real observed population. Note that for the X-ray population (last maps on the right) the top and bottom clusters represent the young magnetar population (with estimated age less than 2 kyr, see Section 5 for more details) and the older XDINS population, respectively.

In order to smooth out abrupt features in the maps due to the random nature of the simulation and stabilize the subsequent machine learning pipeline, we apply a gaussian smoothing kernel. The resulting maps associated with the respective ground truth labels, i.e., the parameter set used to simulate them, will be the input provided to train the simulation-based inference framework explained in Section 4. To further stabilize the training procedure, we standardize maps and labels so that the values have a mean of 0 and a standard deviation of 1. Standardization is performed on a sample basis for the maps and on a dataset basis for the labels (see Graber et al. 2024; Pardo-Araujo et al. 2025, for more details).

4. SIMULATION-BASED INFERENCE

In recent years, the development of machine learning has allowed the emergence of new tools to perform parameter inference. Simulation-based inference (SBI) is a recently developed framework that combines the power of Bayesian statistics and deep learning to optimize and perform parameter inference when using complex model simulators (see Cranmer et al. 2020; Zammit-Mangion et al. 2025, for reviews). In this work, we use a SBI approach called neural posterior estimation (NPE) (e.g. Papamakarios & Murray 2016; Lueckmann et al. 2017; Greenberg et al. 2019; Dax et al. 2021; Mishra-Sharma & Cranmer 2022; Vasist et al. 2023; Barret & Dupourqué 2024), which has been adopted to infer neutron star population parameters in recent population synthesis studies (Graber et al. 2024; Pardo-Araujo et al. 2025; Sautron et al. 2025). We refer the reader to these works for a more complete overview and description of the different SBI methods and of NPE in particular, we only summarize the main points here. In NPE, a neural density estimator is trained to directly map a simulation output, $\boldsymbol{x}$, to the posterior distribution, $\mathcal{P}(\boldsymbol{\theta}|\boldsymbol{x})$. This gives the probability for a set of model parameters, $\boldsymbol{\theta}$, to have generated the output, $\boldsymbol{x}$, through a stochastic simulator. The neural density estimator, q, is parametrised by a neural network, F, with weights, ϕ, i.e., $q_{F(\boldsymbol{x},\boldsymbol{\phi})}$. The network is then optimised by minimizing the following loss function

$$\mathcal{L}(\boldsymbol{\phi}) = -\sum_{i=1}^{N} \ln q_{F(\boldsymbol{x}_i,\boldsymbol{\phi})}(\boldsymbol{\theta}_i) \tag{17}$$

over a training data set $\{\boldsymbol{\theta}_i, \boldsymbol{x}_i\}$ of size N. This loss is minimized when the neural density estimator approximates the true posterior, that is:

$$q_{F(\boldsymbol{x},\boldsymbol{\phi})}(\boldsymbol{\theta}) \approx \mathcal{P}(\boldsymbol{\theta}|\boldsymbol{x}). \tag{18}$$

In this work, we adopt the sequential version of NPE called truncated sequential neural posterior estimation (TSNPE) (Deistler et al. 2022; Pardo-Araujo et al. 2025). The workflow of this algorithm is as follows (Pardo-Araujo et al. 2025):

1. Sample the proposal prior distribution to obtain $\boldsymbol{\theta}_i \sim \mathcal{P}(\boldsymbol{\theta})$ for $i = 1, ..., N$ where N, is the number of simulations we want to use for training.
2. Using the simulator, generate synthetic data $\boldsymbol{x}_i \sim \mathcal{P}(\boldsymbol{x}|\boldsymbol{\theta}_i)$ based on $\boldsymbol{\theta}_i$ from step 1.
3. Train the neural density estimator on the dataset composed of pairs $(\boldsymbol{\theta}_i, \boldsymbol{x}_i)$ obtained in the previous steps.
4. Use the trained neural density estimator to approximate the posterior distribution, $\mathcal{P}(\boldsymbol{\theta}|\boldsymbol{x}_0)$, at the observed data, $\boldsymbol{x}_0$.
5. Restrict the prior distribution to the approximated posterior distribution computed in step 4 (see Section 3.3 in Pardo-Araujo et al. 2025, and references therein).

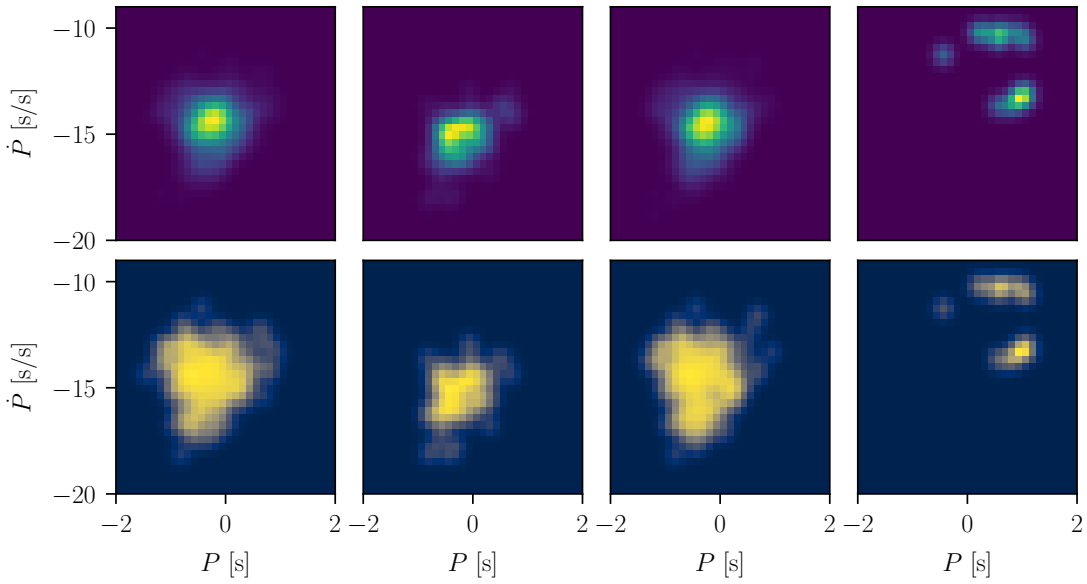

Figure 2. The eight density maps for the observed population of neutron stars. The top and the bottom row show the P-$\dot{P}$ diagrams and the P-$\dot{P}$ averaged flux maps, respectively, for the (from left to right) PMPS, SMPS, HTRU, and X-ray surveys. In the top row, the color represents the density in neutron star number within each bin, while in the bottom row the color represents the averaged flux in Jy for the radio survey maps and in erg s^{-1} cm^{-2} for the X-ray survey map within each bin. For bins without any stars in the bottom row, the averaged flux in log scale has been set to -7 and -17 for the radio and X-ray maps, respectively.

6. Update the proposal prior distribution with the new restricted prior and return to step 1.

Steps 1 to 6 are performed several times and each iteration is commonly referred to as *training round*. This way, the posterior is iteratively refined focusing the computational resources on the region of the parameter space that is more compatible with the observed data. In each round, the newly generated simulations are added to the ones used in the previous rounds, so that the combine simulations are employed together as a training dataset.

4.1. *Neural posterior estimation setup*

In this work, we use the same setup as in Graber et al. (2024); Pardo-Araujo et al. (2025), and only summarize the main aspects important for this study. To set up the inference procedure and the neural density estimator, we use the library `sbi`[3] (v=0.22.0, Tejero-Cantero et al. 2020). The output of the simulations is represented in the form of eight 2D maps as described in Section 3.8 and is processed by a convolutional neural network (CNN) as an eight-channel input. The CNN architecture is constituted by two 2D-convolutional layers, each followed by a 2D max pooling layer and a ReLU activation function (Glorot & Bengio 2010) extracting key features into a latent vector of size 32. This latent vector is the input to the neural density estimator which we set to be a mixture density network (MDN). This MDN is composed of three fully connected layers with 32 neurons each and ReLU activation function, followed by an output layer comprising four fully connected sub-layers, corresponding to the mean, weight, diagonal, and upper triangular components of the covariance matrices for the Gaussian mixture (see Fig. 6 in Graber et al. 2024). We use ten Gaussian components in the mixture to ensure sufficient flexibility when approximating the posterior. We fix the batch size to eight, the fraction of training dataset used for validation to 0.1 and train both the CNN and the MDN simultaneously using the Adam optimiser (Kingma & Ba 2014) with an initial learning rate of 5×10^{-4} and an early stopping criterion of 20 epochs to prevent overfitting. This implies that training is stopped if the validation metric (given by the same Eq. (17) computed over the validation dataset) does not improve for 20 consecutive epochs, with the best validation weights saved. At each training round, the weights for the CNN are re-initialized using the Kaiming prescription (He et al. 2015), while for the MDN the weights are initialised with PyTorch's default initialisation. Moreover, in order to take into account the variability of the initialization and the training process in each round, we train an ensemble of five neural networks with identical architectures as described above. The fi-

[3] https://github.com/sbi-dev/sbi

nal posterior in each round is then given as a weighted sum of the individual posterior distributions computed by each network based on equal weights(see Graber et al. 2024; Pardo-Araujo et al. 2025, for more details). In each of the experiments described below in Sections 5.1 and 5.2, we train the TSNPE algorithm over ten rounds. In the first round, we use 30000 simulations for training and validation and 3000 for testing purposes, and adopt the following uniform prior ranges for the parameters:

$$\begin{aligned}
\mu_{\log P} &\in \mathcal{U}(-1.5, 0.5),\\
\sigma_{\log P} &\in \mathcal{U}(0.1, 1),\\
\mu_{\log B,1} &\in \mathcal{U}(12, 13.5),\\
\sigma_{\log B,1} &\in \mathcal{U}(0.1, 1),\\
\mu_{\log B,2} &\in \mathcal{U}(13.5, 15),\\
\sigma_{\log B,2} &\in \mathcal{U}(0.1, 1),\\
w_{\log B} &\in \mathcal{U}(0.1, 1),\\
a_{\rm late} &\in \mathcal{U}(-3, -0.5),\\
\mu_{\log L_0} &\in \mathcal{U}(24.6, 28.6),\\
\alpha_L &\in \mathcal{U}(0.1, 1).
\end{aligned} \tag{19}$$

In each of the following rounds, we generate 1000 new simulations by sampling the restricted prior. This number has been chosen arbitrarily but as in Pardo-Araujo et al. (2025), it guarantees a good compromise between computational costs and network convergence as shown below. We then retrain the ensemble on the new dataset formed by adding the new simulations to the ones of the previous rounds. The training process is executed on a Tesla V100 SXM2 GPU with 32 GB of memory. The generation of simulations in each TSNPE round for both the training and test datasets are parallelized to speed up the algorithm. For this, we use the Python package Dask (Dask Development Team 2016), a library for dynamic task scheduling. In total, 600 CPU workers are employed to handle the parallelized simulations.

5. RESULTS

5.1. *Inference on the entire population*

We first train the TSNPE algorithm over 10 rounds with the setup explained in Section 4.1 considering the entire observed population of radio and X-ray emitting neutron stars described in Section 2.

After round 5 the approximated posterior results are stable therefore we use the model trained in this round to infer the best parameter values. In the first column of the table 2 we list the median of the best inferred parameters from round 5 with the 95% credible interval.

We sample 100 parameter sets from this posterior distribution and generate the respective simulations. In this way we take into account the uncertainty on the inferred parameters and their correlation. We then compare the results of these simulated populations with the real observed populations. In Fig. 3 we show the comparison between our best simulations and the observed dataset in the $P - \dot{P}$ diagram. The contour lines represent the number density of the simulated populations plotted over the real observed radio pulsars (dots) and X-ray emitting neutron stars (crosses). Both the radio pulsar population as well the entire X-ray emitting neutron star populations are well reproduced by the simulations.

However if we focus our attention on the sub-populations of young magnetars (pink crosses and contour lines) and of the XDINSs (orange crosses and contour lines), our simulations struggle to reproduce the observations. To define the sub-population of young magnetars we follow Pardo-Araujo et al. (2026) and select from the observed dataset the ones whose age estimated from a supernova remnant association or from the characteristic age is lower than 2 kyr (pink crosses). To compare with this sub-sample in our simulations we then select the neutron stars with simulated ages lower than 2 kyr (pink contour lines). For the sub-population of XDINS-like sources (orange contour lines) we consider the simulated neutron stars within a distance of ~ 500 pc from the Sun and an age greater than 10^5 yr to mimic the properties of the observed XDINSs. While the simulations are able to predict a number of detected young magnetars of 6 ± 3, which is compatible with the seven observed, they predict their distribution to be shifted at lower spin periods and spin-period derivatives. On the

Table 2. Comparison between best parameters values inferred by considering the whole sample of radio pulsars and X-ray emitting neutron stars (Radio + X-ray) and by considering the radio pulsars and a restricted sample with only the young magnetars and the XDINSs (Radio + young mag, XDINSs).

Parameter	Radio + X-ray	Radio + young mag, XDINSs
$\mu_{\log P}$	$-0.19^{+0.45}_{-0.73}$	$-0.26^{+0.44}_{-0.46}$
$\sigma_{\log P}$	$0.61^{+0.28}_{-0.33}$	$0.70^{+0.23}_{-0.34}$
$\mu_{\log B,1}$	$12.70^{+0.26}_{-0.23}$	$12.72^{+0.20}_{-0.21}$
$\sigma_{\log B,1}$	$0.45^{+0.11}_{-0.12}$	$0.45^{+0.13}_{-0.13}$
$\mu_{\log B,2}$	$13.85^{+0.41}_{-0.32}$	$14.11^{+0.53}_{-0.51}$
$\sigma_{\log B,2}$	$0.35^{+0.18}_{-0.19}$	$0.49^{+0.35}_{-0.32}$
$w_{\log B}$	$0.83^{+0.16}_{-0.22}$	$0.74^{+0.21}_{-0.25}$
$a_{\rm late}$	$-0.89^{+0.31}_{-0.32}$	$-0.85^{+0.26}_{-0.26}$
$\mu_{\log L_0}$	$25.78^{+0.24}_{-0.24}$	$25.70^{+0.25}_{-0.26}$
α_L	$0.71^{+0.09}_{-0.09}$	$0.73^{+0.09}_{-0.09}$

other hand, for the XDINSs the simulations are able to reproduce their timing properties but they under-estimate their number to be 2 ± 2 compared to the seven observed sources.

By counting the number of neutron stars we need to generate in each simulation to reach the observed detected numbers in each survey (see Section 2) we can estimate a neutron star birth rate for each of the simulated survey. The estimated mean and standard deviation of the birth rates for the different surveys are the following:

$$\begin{aligned} &\text{PMPS: } \sim 3.6 \pm 0.8 \text{ neutron stars per century,} \\ &\text{SMPS: } \sim 3.0 \pm 0.7 \text{ neutron stars per century,} \\ &\text{HTRU: } \sim 3.0 \pm 0.7 \text{ neutron stars per century,} \\ &\text{X-ray: } \sim 3.1 \pm 0.9 \text{ neutron stars per century,} \end{aligned} \tag{20}$$

Note that these values are compatible between each other at 1-σ level but are higher than the core collapse supernovae rate of 1.63 ± 0.46 per century estimated by Rozwadowska et al. (2021).

In Fig. 4 we compare the observed $\log N - \log S$ distribution for the absorbed X-ray fluxes (black line) with the one predicted by the simulations (purple line and bands). The observed $\log N - \log S$ is marginally consistent falling inside the 3-σ uncertainty of the simulations. However, in our simulations we observe a lack of bright sources with fluxes above $\sim 10^{-12}\,\mathrm{erg\,s^{-1}\,cm^{-2}}$. Since the tail at high fluxes is where young or nearby bright sources dominate and where the completeness level of the observed sample is higher, this lack of sources in our model is another indication that the simulations either under-estimate the detected numbers or the observed fluxes for these two sub-populations (see Section 6 for more discussion about this aspect).

5.2. *Inference on the sub-sample of young X-ray population and XDINSs*

As a second experiment we check the impact on our inference results of considering only the restricted samples of young magnetars and XDINSs as defined in the previous section. This should represent a less biased and more complete sample. Young sources are intrinsically rare and X-ray bright, therefore, observational biases are not expected to significantly affect their detection. XDINSs on the other hand are observed to have a high X-ray flux due to their proximity. These should represent an almost complete sample considering a limiting distance of $\sim$ 500 pc from the Sun (in Kurpas et al. 2026, they discuss possible new candidates, with only one being closer than 500 pc). This results in a total of 13 X-ray bright sources to use for inference beside the radio pulsar population. We generate the maps using this extra filter on the X-ray population for both the simulations and the observed sample and train the TSNPE algorithm. The inferred posterior distribution conditioned on these restricted observed data for each round are displayed in Fig. 5. As in the previous experiment, overall the posterior distribution is stable over the training rounds.

In order to compare the simulated population from the best parameters inferred from the posterior with the real one, we choose to use the posterior estimate from round 5 as it displays conservative results which encompass the ranges of the posterior estimated in the following rounds. To assess the quality of the trained posterior estimator in this round we check its predictive power over the test dataset by comparing its posterior predictions compared to the ground truth labels (see Fig. 11). Moreover we computed the coverage probability over the same test dataset (see Fig. 12) which shows that the approximated posteriors are conservative since the coverage probability is overall greater than the credibility level (Hermans et al. 2021; Graber et al. 2024; Pardo-Araujo et al. 2025, see Appendix B for more details). These aspects highlight the robustness of the inference results. In the second column of table 2 we list the median of the best inferred parameters from round 5 with the 95% credible interval.

As in the previous section we sample 100 parameter sets from this posterior distribution and generate the respective simulations. We then compare the distributions of these simulated populations with the real observed populations in the $P - \dot{P}$ diagram (see Fig. 7). The radio pulsar population is still well reproduced by the simulations. On the other hand, when looking at the X-ray population, while this time the young magnetars (pink crosses) and XDINSs (orange crosses) are well reproduced by the simulations, the bulk of the mock observed X-ray emitting neutron stars (purple contour lines) is shifted towards longer periods when compared with the real data.

The estimated mean and standard deviation of the birth rates for the different surveys are the following:

$$\begin{aligned} &\text{PMPS: } \sim 3.6 \pm 0.9 \text{ neutron stars per century,} \\ &\text{SMPS: } \sim 3.0 \pm 0.8 \text{ neutron stars per century,} \\ &\text{HTRU: } \sim 3.0 \pm 0.7 \text{ neutron stars per century,} \\ &\text{X-ray: } \sim 1.3 \pm 0.9 \text{ neutron stars per century,} \end{aligned} \tag{21}$$

Note that also in this case the predicted core collapse supernovae rate is larger than the value from Rozwadowska et al. (2021). In particular while the values of the three radio surveys are still compatible between each other

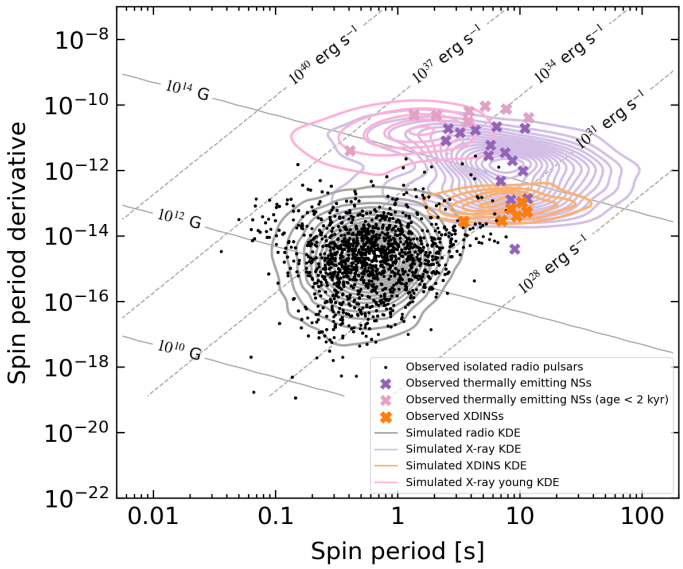

Figure 3. Comparison in the $P - \dot{P}$ diagram between the observed population of radio pulsars and X-ray emitting neutron stars and the best simulated populations assuming a double log-normal initial magnetic field distribution. Points and crosses represent the observed radio pulsars and X-ray emitting neutron stars. Pink crosses represent the young (estimated ages lower than 2 kyrs) magnetars, while orange crosses the observed XDINSs. The contour lines represent density contours obtained from 100 simulations generated from sets of input parameters drawn from the posterior distribution obtained from round 5.

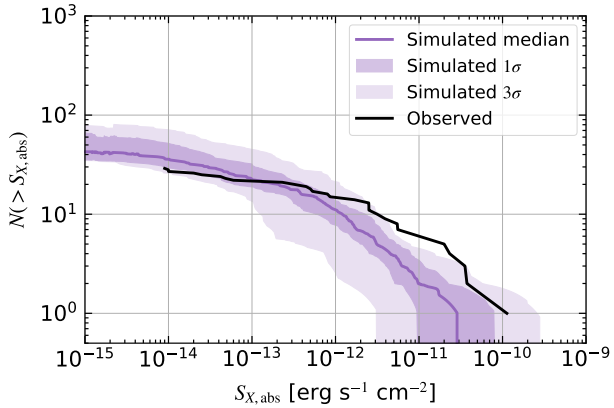

Figure 4. Comparison of the $\log N - \log S$ distributions for the observed and simulated X-ray thermally emitting neutron stars. The black line represents the trend for the observed population while the purple line and bands represent the median, the 1-σ and 3-σ uncertainties computed over 100 simulations generated from sets of parameters drawn from the posterior distribution obtained from round 5.

at 1-σ level, for the X-ray survey the mean birth rate is significantly lower compared to the others. This is an indication that this time our simulations are overestimating the number of detected X-ray emitting neutron stars compared to reality. However while overestimating the numbers for the entire X-ray detected population, these simulations predict 8 ± 5 young magnetars and 4 ± 3 XDINS-like sources allowing a better agreement with the real observed numbers, i.e., 7 and 7 respectively.

In Fig. 8 we compare the $\log N - \log S$ distributions. In the left panel we show the distribution when considering only the young magnetars and XDINS-like sources. In this case, the observed $\log N - \log S$ falls inside the 1-σ uncertainty of the simulations. When considering the whole X-ray population (right panel) the simulations are compatible at 1-σ level with the high-flux tail of the observed distribution. On the other hand at lower fluxes, while being still compatible with the observed distribution at a 3-σ level, the simulations predict more sources at flux values between $10^{-14} - 10^{-12}$ erg s^{-1} cm^{-2}. This excess is also reflected in the $P - \dot{P}$ diagram where the number of mock detected sources with fluxes in this range but with spin periods in the range between 10 to 100 s is overestimated, as highlighted above. Never-

theless, it is important to note that the simulations are able to reproduce the tail at high fluxes where young or nearby bright sources dominate and where the completeness level of the observed sample is higher.

In Fig. 9 we show the best fit distribution for the initial magnetic fields together with its uncertainty. In particular we draw 10000 values from the posterior distribution obtained from round 5 (see Fig. 6) and evaluate the median distribution (purples solid lines) and the 1σ and 3σ uncertainties (dashed regions).

6. DISCUSSION

Compared with the results on radio pulsars obtained in Pardo-Araujo et al. (2025), here we find that the initial spin period distribution is shifted to larger values even if the values are consistent within the uncertainties. The uncertainty on the parameters describing this distribution is still large and indicates that these parameters are difficult to constrain as already discussed in previous works. The memory of the initial spin periods is rapidly lost during the magneto-rotational evolution. Adding the sample of X-ray emitting neutron stars at higher magnetic field values does not help in these regards as the stronger the magnetic field the faster is the evolution in the $P-\dot{P}$ diagram.

As $\mu_{\log P}$ is negatively correlated with the initial magnetic field parameters (see Fig. 6), an initial spin period distribution shifted to larger values allows for lower values of $\mu_{\log B,1}$. Therefore the first gaussian component of the magnetic field distribution is shifted to slightly lower values compared to Graber et al. (2024); Pardo-Araujo et al. (2025). This is also compensated by the addition of the second component at stronger magnetic fields that helps providing sources with magnetic fields above 10^{13} G that can be detected as radio pulsars (see Fig. 9).

As in Pardo-Araujo et al. (2025) we find a slightly negative correlation between the radio luminosity parameter, $\mu_{\log L_0}$, and $\mu_{\log P}$, and a slightly positive correlation between $\mu_{\log L_0}$ and $\mu_{\log B,1}$. This means that as the radio pulsar population shifts to the right and to the bottom of the $P-\dot{P}$ diagram because of a larger $\mu_{\log P}$ and a smaller $\mu_{\log B,1}$ respectively, the radio luminosity can also be slightly smaller as otherwise too many pulsars with low rotational energy would be detected.

Turning to the parameters relevant for the X-ray population, we find that when inferring on the entire sample of X-ray emitting neutron stars, the second component of the initial magnetic field is centered at values of $\sim 7\times10^{13}$ G. On the other hand when focusing the inference on the sub-populations of young magnetars and XDINSs the mean value of this component shifts to $\sim 1\times10^{14}$ G. This difference is due to the fact that in the first case, when fitting the entire observed X-ray population, the TSNPE algorithm tries to reproduce the spin period distribution which shows a clear cut off at periods of around $\sim 10-20$ s. To achieve this, the initial magnetic field cannot be too large, otherwise, the spin-down would be too strong, leading to final spin periods larger than those observed. However, by limiting the initial magnetic field strength the simulations poorly reproduce the high $\dot{P}$ values of the young magnetars as can be seen by looking at the pink contour lines and crosses in Fig. 3. Moreover, another counter effect is the reduction of the X-ray luminosities and the observed fluxes, which is indeed evident in the $\log N-\log S$ distribution (see Fig. 4). On the other hand, when focusing the inference on the young magnetars and XDINSs, a second component with stronger initial magnetic fields can reproduce both sub-populations. This implies also higher thermal luminosities and a better fit to the high-flux tail of the $\log N-\log S$ distribution (see Fig. 8). However, the resulting population of X-ray emitting neutron stars tends to have spin periods longer than those observed due to a stronger spin-down (see Fig. 7).

In general the uncertainty on the values of the parameters describing the second component of the initial magnetic field distribution is large as can be observed by the large variance associated to this component in Fig. 9. This is a consequence of the low number statistics for the X-ray emitting neutron stars (see also Appendix B). This can indeed introduce random statistical noise in the maps of the simulated and observed samples which makes the training process more challenging and the inference results less precise. On the other hand this could be also an indication that the chosen parametrisation describing the shape of the initial magnetic field distribution or the adopted models for the evolution of magnetar-like sources in the $P-\dot{P}$ diagram or the observational biases are not describing faithfully the reality.

6.1. *The ~ 20 s cut-off in the magnetar population*

The presence of a cut-off in the observed spin period distribution at around ~ 20 s is intriguing and it has been challenging to explain in previous population synthesis studies (see for example Gullón et al. 2015). The fact that also our models struggle to reproduce this sharp feature could be due to different aspects.

First of all our treatment of the outburst rate is simplified, primarily as we rely on 2D magneto-thermal simulations where failures are obtained under simplified assumptions (see Dehman et al. 2020, for more details), secondly as we consider that all crust failures events lead to a detectable outburst. This picture does not

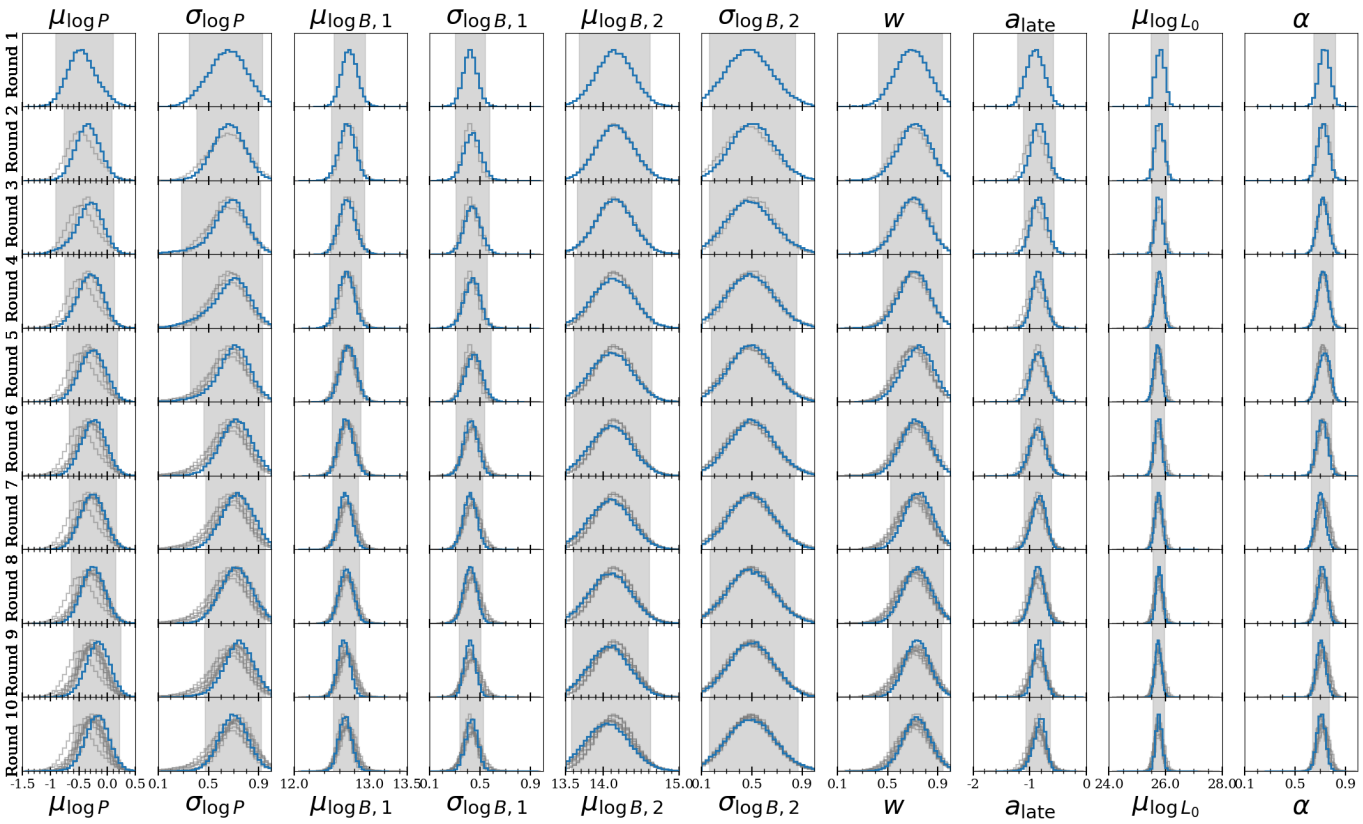

Figure 5. Results of the TSNPE algorithm applied to our population synthesis where we assumed a double log-normal initial magnetic field distribution. Each row shows the marginalized posterior obtained in each round of inference where the columns refers to different parameters. The last column shows the coverage probability computed on the test dataset. This is only computed in the first round. In each panel, for a given parameter the marginalized posterior computed in that specific round id shown in blue, while the marginalized posteriors from previous rounds are shown in light gray. The gray shaded area represents the 95% credibility interval of the approximated posterior for that round. In each of these 1-D marginal posterior distribution, the horizontal axes represent the parameters' prior ranges.

take into account the energy of these events and their luminosity in gamma-ray and X-ray photons. Quantifying the efficiency with which the energy dissipated by the stresses is converted into photon luminosity that could be detected by gamma-ray and X-ray instruments is challenging (Lander 2023; Qu & Bransgrove 2026). Moreover we are not taking into account sensitivities and limitations of current and past X-ray observatories for transient detection. One issue for example is that outbursts occurring during periods when the source is Sun-constrained, i.e., too close to the Sun to be observed, will be missed. This happens for a few months each year for every source effectively reducing the detectability time-window; Therefore, combining all these aspects, in our simulations we are likely overestimating the number of sources that could be detectable due to an outburst event and this could also lead to an over-detection of sources at long spin periods which have not been observed in reality.

A second aspect to consider is that our results rely on the assumption that the configuration of the initial magnetic field has only a dipolar poloidal and toroidal components. For the magnetar population this could not hold, the presence of strong multipole components beside the dipole (Dehman et al. 2023b, e.g.) could affect how the magnetic field evolves, implying a faster decay and an enhanced dissipation in the neutron star crust that could lead to higher X-ray luminosities in the first stages of the magnetar evolution. A faster decay could help explain the sharp drop of sources at a spin period of ~ 20 s as the spin-down evolutionary trajectory would bend down sooner in the $P - \dot{P}$ diagram. A different configuration of the field also affects the rate of expected crustal failures and consequently of outburst events. The presence of strong multipolar components in general should enhance the stresses in the crust and the rate of failures leading to more outburst events and a greater number of detectable magnetars.

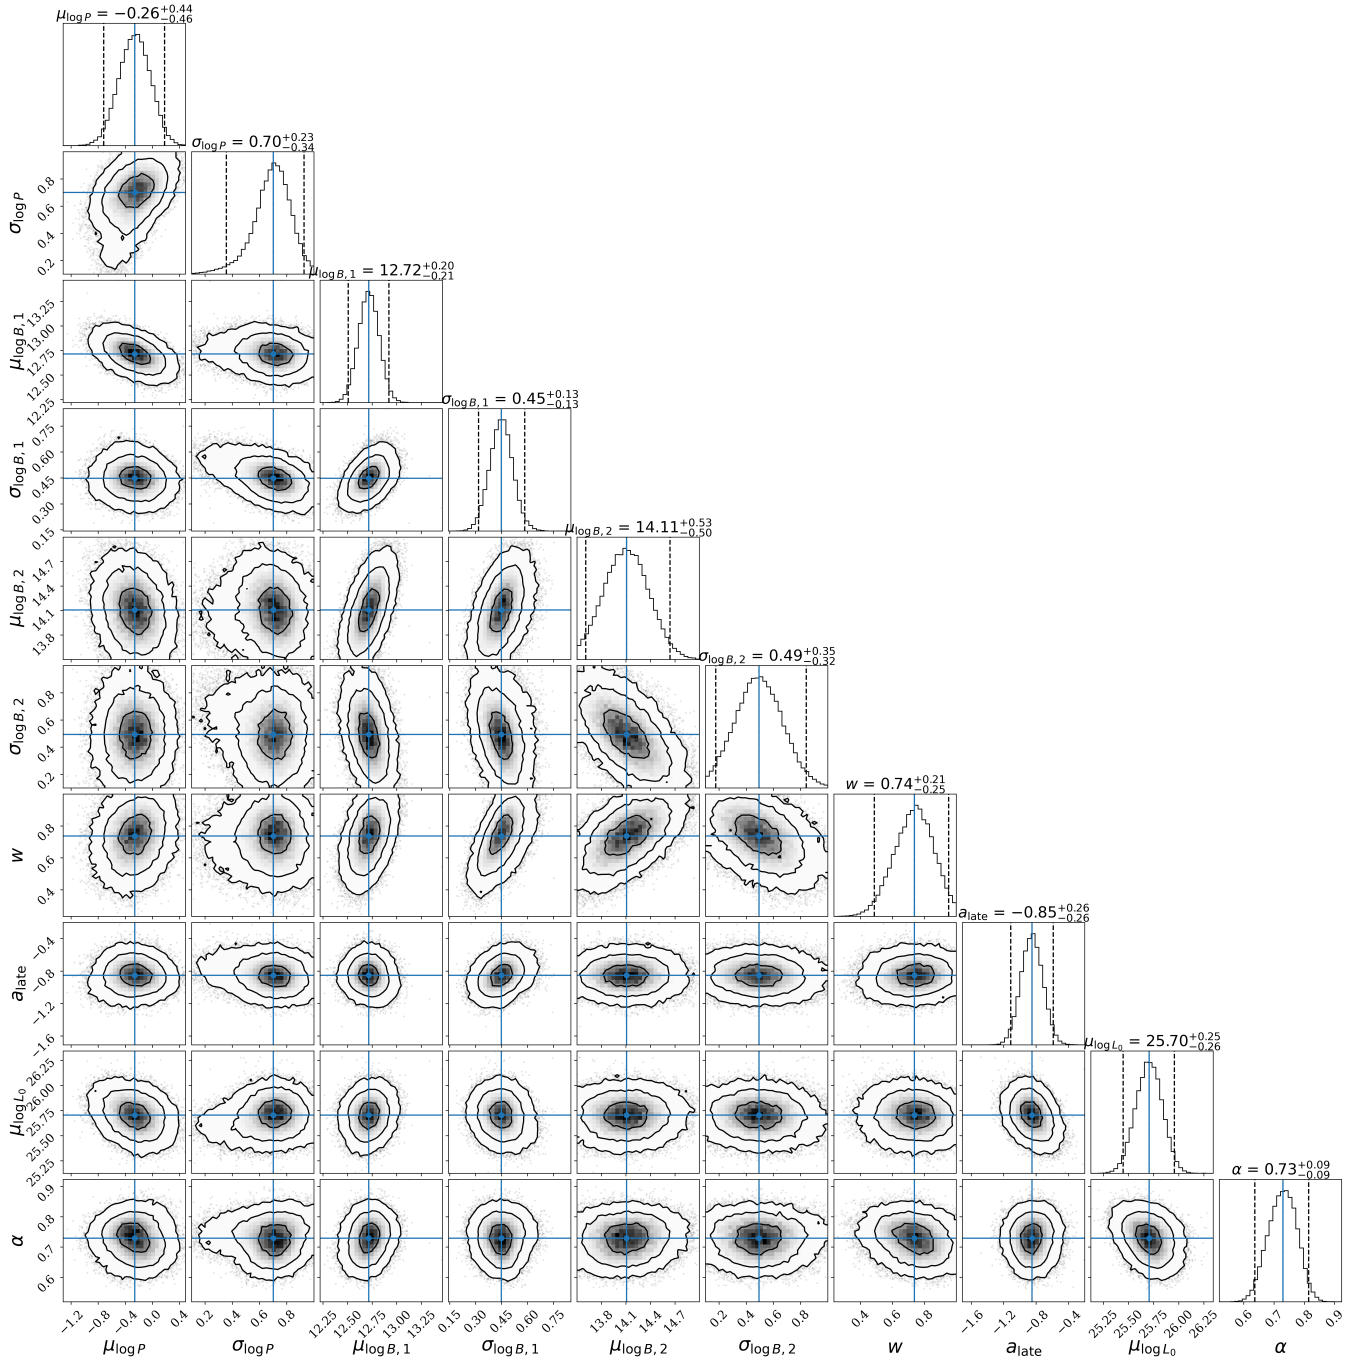

Figure 6. Corner plot of the posterior distribution conditioned on the observed data restricted to the radio pulsars and the sub-populations of young magnetars and XDINSs. The posterior is estimated from the round 5 of the TSNPE algorithm and adopted for the best parameter estimation in Section 5.2.

A second aspect could be related to the evolutionary models. In the case of young magnetars their frequent flares and outburst activity could cause changes in the magnetosphere configuration and particle winds that could momentarily enhance the spin-down rate (Tong et al. 2013; Pétri 2019). In this case assuming a force-free dipolar magnetospheric spin-down model could be an over-simplification and could bias the estimate of the surface dipolar magnetic field of young magnetars and therefore their birth properties. A different evolutionary track in the $P-\dot{P}$ diagram where the $\dot{P}$ is enhanced could also contribute to explain the timing properties of magnetars. However an enhanced spin-down acting for long time would lead to longer final spin periods which would be against the observed cut-off at ~ 20 s. If the enhanced spin-down is only temporary this would imply

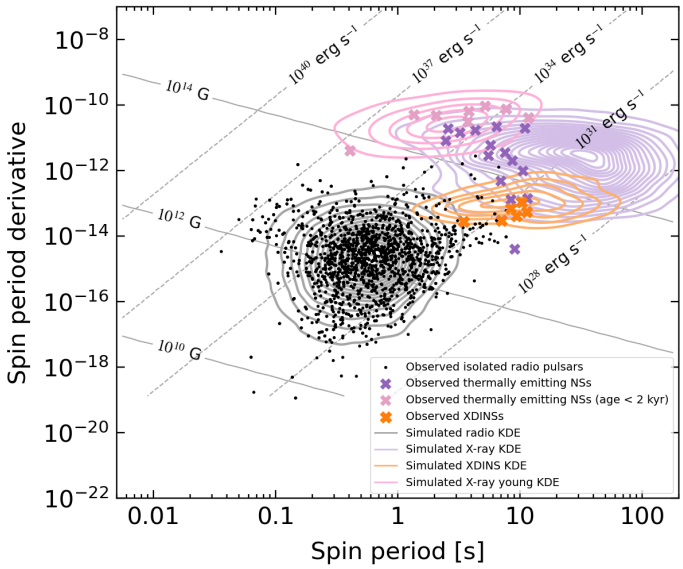

Figure 7. Comparison in the $P - \dot{P}$ diagram between the observed population of radio pulsars and X-ray emitting neutron stars and the best simulated populations assuming a double log-normal initial magnetic field distribution. Points and crosses represent the observed radio pulsars and X-ray emitting neutron stars. Pink crosses represents the young (estimated ages lower than 2 kyrs) magnetars, while orange crosses the observed XDINSs. The contour lines represent density contours obtained from 100 simulations generated from sets of input parameters drawn from the posterior distribution obtained from round 5.

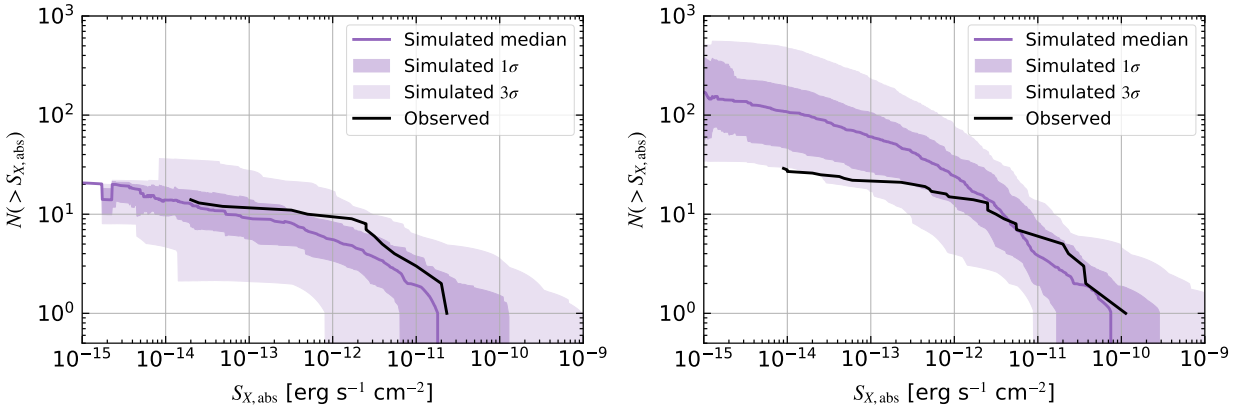

Figure 8. Comparison of the $\log N - \log S$ distributions for the observed and simulated X-ray thermally emitting neutron stars assuming a double log-normal initial magnetic field distribution. The left panel considers only the young magnetars (with age lower than 2 kyr) and XDINS-like sources. The right panel represents the entire observed X-ray population. In both plots the black line represents the trend for the observed population while the purple line and bands represent the median, the 1-σ and 3-σ uncertainties computed over 100 simulations generated from sets of parameters drawn from the posterior distribution obtained from round 5.

that we are detecting magnetars only in this short temporary phase which would make such detections more unlikely unless a strong detection biases like an outburst event is present.

6.2. *Estimate of the magnetar fraction*

When inferring considering the entire X-ray population we find that the second component of the initial magnetic field distribution has a relative weight of 17^{+22}_{-16}%. On the other hand when focusing on the young magnetars and XDINSs we find 26^{+25}_{-21}%. This number can also be interpreted as the fraction of neutron stars born with a strong magnetar-like magnetic field. These values are consistent within the uncertainties with the value $10.7^{+18.8}_{-4.4}$% inferred by Sautron et al. (2025) who performed a simulation based inference study although only considering magnetars and with the estimates of Pardo-Araujo et al. (2026) who put constraints on the fraction of neutron stars born as magnetars by combining the young population of neutron stars and their supernova association, CCOs and XDINSs. While using the same simulation framework and similar assumptions

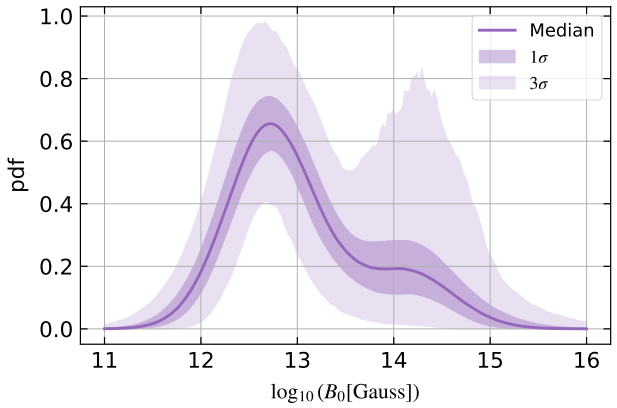

Figure 9. Best fit distribution of initial magnetic fields for the inference experiment on radio pulsars, young magnetars and XDINS. The purple solid line represents the median value while the shaded regions represent the 1σ and 3σ uncertainties computed over 10000 random samples of the posterior distribution shown in Fig. 6.

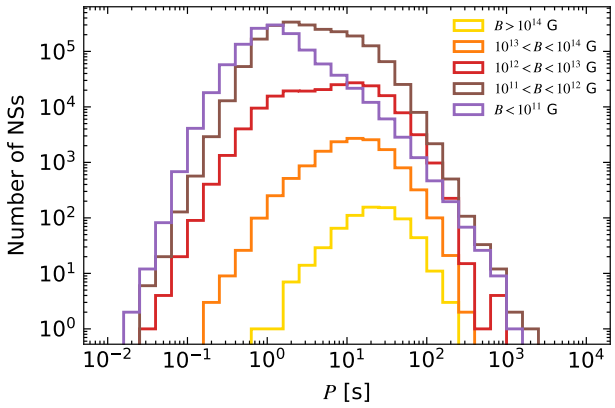

Figure 10. Distribution of the final spin period of the entire evolved neutron star population as a function of different ranges of final magnetic fields.

as in Pardo-Araujo et al. (2026) the approach here is slightly different as we focus on the entire evolved population. This makes this study and results more sensitive to the modelling of the correct observational biases and evolutionary scenarios. In Pardo-Araujo et al. (2026), we found a conservative range for the magnetar fraction of 30–70%. Their Extended Data Figure 4 shows that the inferred magnetar fraction decreases as the neutron star birth rate increases. Therefore, our lower estimated magnetar fraction is consistent with the higher birth rates we predict (3–4 neutron stars per century, see Eq. 21). Overall if we combine these estimates we can argue that most likely the Galactic neutron star birth rate lies in the range 3-4 neutron stars per century of which a fraction of $\sim 30-50\%$ are born as magnetars. Constraints on the magnetar fraction are important to assess the impact of these progenitors as possible engines of extreme events such as FRBs (Beniamini et al. 2023, 2025), superluminous supernovae (Inserra et al. 2013), as well as the plateau phases observed in many GRB afterglows (Metzger et al. 2011; Lü & Zhang 2014).

6.3. *The long-period pulsar subpopulation*

Recently a pulsar with a spin period of ~ 76 s (Caleb et al. 2022) and several Long Period Transients (LPTs), recurring on timescales ranging from hundreds of seconds up to several hours, have been discovered in radio (see Rea et al. 2026, for a review). Neutron stars and in particular magnetars have been suggested as possible engines to produce the emission of some of these LPTs. In the assumptions of our best models, in order to see how many neutron stars we expect in our Galaxy to have ultra-long spin period we performed a simulation adopting the best parameter values inferred in Section 5.2. To simulate the full population we assume a birth rate of 3.5 neutron stars per century and evolve it up to a maximum age of 10^8 yrs for a total of 3.5×10^6 neutron stars. Note that the birth rate value we consider here is consistent with the values obtained in Section 5.2. In Fig. 10 we report histograms of the final spin periods considering different ranges of final magnetic fields for the entire population. We predict that the total number of sources with spin periods above 20 s should be $\sim 2 \times 10^5$ sources ($\sim 6\%$ of the entire population) with ~ 400 objects ($\sim 0.01\%$) with still magnetar-like magnetic fields above 10^{14} G. In particular we expect the number of neutron stars with spin periods above 100 s to be ~ 5000 sources ($\sim 0.1\%$). We note that in our simulation framework neutron stars with spin periods longer than a few seconds becomes hard to detect. This is due to a combination of factors. First of all we assume a radio beam aperture which has an inverse dependence on the spin period as $\propto P^{-1/2}$ (see Eq. 13 in Graber et al. 2024, and references therein). Moreover neutron star reaching long spin periods are usually the ones with stronger magnetic fields and therefore a stronger tendency of evolving towards being aligned rotators, i.e. with magnetic field dipole axis aligned with the rotation axis (see Eq.(4)). These aspects naturally reduce the probability of the radio beam intersecting our line of sight for slow rotators. Furthermore without considering alternative evolutionary spin-down models (see, e.g., Ronchi et al. 2022; Gençali & Ertan 2024; Zhou et al. 2024) and assuming dipolar spin-down only, we predict that very few sources could reach spin periods around 1000 s. Therefore without other mechanism to enhance spin-down our models could explain only the fast spinning end of the LPT population.

6.4. *The source populations of FRBs*

There is broad and convincing evidence that FRBs are emitted by young, energetic neutron stars (e.g., Pastor-Marazuela et al. 2025). As the bursts are, on average, 10^{12} times brighter than normal pulsars (cf. van Leeuwen et al. 2023), the inferred energy budget suggests that magnetars, specifically, could be emitting these bursts. The case is strengthened by both single-object and population evidence: the detection of a short bright burst from magnetar SGR 1935+2154 (Bochenek et al. 2020), the agreement in burst wait-time statistics between specific repeaters and magnetars (e.g. Wadiasingh & Timokhin 2019), and the finding that the number density, spectral index and dispersion measures of the population of FRBs agrees qualitatively with the expectations for a magnetar origin (Gardenier & van Leeuwen 2021). The population synthesis method we described here has the potential to make this link much more quantitative. By connecting our radio and X-ray results to the state-of-the-art in FRB population synthesis codes (e.g. Wang & van Leeuwen 2024) we will be able to determine, for example, if FRB repetition frequency, too, is guided by the same outburst prescriptions that drive magnetar flares (Sect. 3.5, and bottom panel of Fig. 1). Such a detection would offer strong evidence that FRBs are, indeed, distant magnetars.

7. CONCLUSIONS

We performed a simulation based inference study to constrain the Galactic population of neutron stars. Our approach was unique in constructing a unified framework for the populations of radio pulsars, magnetars, and XDINSs with a significant thermal component in their spectra. To achieve this we developed `ML-Poppyns`, a population synthesis framework that combines models to simulate the birth, evolution and detection of a Galactic population of neutron stars, with simulation based inference that exploit the power of neural networks to perform parameter inference. To reproduce the properties of magnetars and XDINSs, we incorporate results from state-of-the-art 2D magneto-thermal simulations to model the coupled evolution of magnetic fields, surface temperature and crustal failures.

Using these, we determine which source population best describes the observed one. This underlying population can help refine the expected detection numbers for future instruments like the Square Kilometer Array (SKA Keane et al. 2015, 2025) in the radio band and eXTP (Santangelo et al. 2019) and NewAthena (Cruise et al. 2025) in the X-ray band. It also can serve as a starting point for investigating the impact of strongly magnetic neutron stars as central engines for powerful events such as FRBs and GRBs.

We show that an initial magnetic field distribution described by two log-normal components with means of 5×10^{12} G and 1×10^{14} G respectively, is able to reproduce the properties of radio pulsars, young magnetars and XDINSs. It does, however, struggle to explain the absence of detected sources with spin periods larger than ~ 20 s. We also find that to simultaneously explain the detected numbers at both electromagnetic wavelengths, we require a Galactic neutron star birth rate of 3-4 neutron stars per century, with a 10-50% magnetar fraction.

M.R. and J.v.L. are supported by the Dutch Research Council (NWO) via the grant CORTEX (NWA.1160.18.316) of the research programme NWA-ORC. V.G. is supported by a UKRI Future Leaders Fellowship (grant number MR/Y018257/1). C.P.A. and N.R. are supported by the ERC via the Consolidator grant "MAGNESIA" (No. 817661), the ERC Proof of Concept "DeepSpacePULSE" (No. 101189496), and by the program Unidad de Excelencia María de Maeztu CEX2020-001058-M. We also acknowledge support from the Catalan grant SGR2021-01269 (PI: Graber/Rea) and the Spanish grant PID2023-153099NA-I00 (PI: Coti Zelati). C.P.A.'s work has been carried out within the framework of the doctoral program in Physics at the Universitat Autonoma de Barcelona. C.D. is supported by the Ministerio de Ciencia, Innovación y Universidades (JDC2023-052227-I), co-funded by AEI (MCIN/AEI/10.13039/501100011033), the FSE+, and the Universidad de Alicante. C.D. and J.P acknowledge support from the Conselleria d'Educació, Cultura, Universitats i Ocupació de la Generalitat Valenciana (grant CIPROM/2022/13). D.D.G. is supported by a Juan de la Cierva fellowship (JDC2023-052264-I). AM acknowledges support from the Fund Vera Rubin/Chile 2024, under the project DIA 1736 "Silent black holes around red supergiants". F.C.Z. is supported by a Ramón y Cajal fellowship (grant agreement RYC2021-030888-I).

The data production, processing, and analysis tools for this paper have been implemented and operated at the Port d'Informació Científica (PIC) data center. PIC is maintained through a collaboration of the Institut de Física d'Altes Energies (IFAE) and the Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas (Ciemat). We particularly thank Christian Neissner and Martin Børstad Eriksen for their support at PIC.

We made use of the pulsar population synthesis code ML-Poppyns Ronchi et al. (2021); Graber et al. (2024); Pardo-Araujo et al. (2025) funded by the European Research Council via the ERC Consolidator grant "MAGNESIA" (No. 817661; PI: N. Rea), and publicly available at https://ice-csic-astroexotic.github.io/code/ml_poppyns/.

The authors thank Roberto Turolla for useful exchanges on the X-ray emission and resonant cyclotron scattering and Emilie Parent for useful discussions on the radio detection.

Software: Astropy (Astropy Collaboration et al. 2013, 2018), healpy (Górski et al. 2005; Zonca et al. 2019), IPython (Perez & Granger 2007), JupyterLab, Matplotlib (Hunter 2007), Numba (Lam et al. 2015), NumPy (Oliphant 2006; van der Walt et al. 2011; Harris et al. 2020), Pandas (McKinney 2010), PyGEDM, PyTorch (Paszke et al. 2019), sbi (Tejero-Cantero et al. 2020), SciPy (Jones et al. 2001; Virtanen et al. 2020), Sphinx.

APPENDIX

A. RESONANT CYCLOTRON SCATTERING

The non-relativistic classical resonant scattering cross section can be approximated as (see Canuto et al. 1971; Nobili et al. 2008; Yamasaki et al. 2020):

$$\sigma_{\rm res}(\omega) \sim \pi^2 \frac{e^2}{m_e c}\delta(\omega-\omega_B)(1+\cos^2\theta), \tag{A1}$$

where θ is the angle between the photon trajectory and the direction of the magnetic field line.

The effectiveness of the process is quantified by the resonant scattering optical depth $\tau_{\rm res}$ which can be estimated by the integral:

$$\begin{aligned}\tau_{\rm res} &= \int n_e \sigma_{\rm res}(\omega){\rm d}l \\ &\sim \int_{R_{\rm NS}}^{r} n_e \sigma_{\rm res}(\omega){\rm d}r' \\ &= \frac{\pi^2 e^2 n_e}{m_e c}(1+\cos^2\theta)\int_{R_{\rm NS}}^{r}\delta(\omega-\omega_B(r')){\rm d}r' \\ &= \frac{\pi^2 e^2 n_e}{m_e c}(1+\cos^2\theta)\frac{r}{3\omega_B} \\ &= \tau_0(1+\cos^2\theta),\end{aligned} \tag{A2}$$

where the integral is performed along the line of sight l, for simplicity taken to be the radial distance r from the star surface, n_e is the density of the magnetospheric plasma (assumed to be constant), and τ_0:

$$\tau_0 = \frac{\pi^2 e^2 n_e r}{3 m_e c \omega_B}. \tag{A3}$$

Note that ω_B depends on the distance r since in the simplifying assumption of a poloidal dipole field we have that $B(r) \sim B_{\rm P} R_{\rm NS}^3/r^3$ where $B_{\rm P}$ is the magnetic field strength at the pole. Therefore to solve the integral we used the property of the Dirac delta function:

$$\int_{-\infty}^{+\infty}\delta(f(x)){\rm d}x = \sum_i \frac{1}{|f'(x_i)|}, \tag{A4}$$

where x_i are the zeros of the function $f(x)$. In our case $f(x)$ is $\omega-\omega_B(r)$ and its derivative with respect to r is $3\omega_B/r$. The zeros of our function are the values of the radial distance r where $\omega_B(r)=\omega$, i.e., where the frequency of the photon matches the cyclotron frequency.

To compute the spectrum distorted by resonant cyclotron scattering we used the simplified semi-analytical 1D model described in Lyutikov & Gavriil (2006) (see also Rea et al. 2008). In this model it is assumed that the seed thermal photons coming from the surface propagate in the radial direction. They interact with the magnetospheric electron-positron plasma where particles are gyrating along the magnetic field lines and are assumed to have a top-hat thermal velocity distribution centered at zero and extending up to velocities $\pm\beta_{\rm T}$. For simplicity it is assumed that photons propagate parallel to the magnetic field lines, either away or towards the star, hence in Eq.(A2), $\cos\theta = \pm 1$ and $\tau_{\rm res} = 2\tau_0$. A photon with a given energy $E_0 = \hbar\omega_0$ will resonantly interact with electrons that are gyrating with a cyclotron frequency that matches their energy. This will define the distance from the star r where the interaction will happen. The probabilities that a photon with energy $E_0 = \hbar\omega_0$ undergoes resonant scattering and is transmitted or reflected are given by Eqs.(35) in Lyutikov & Gavriil (2006) which we report here:

$$\begin{aligned}p_+(\eta){\rm d}\eta &= e^{-\tau_0/2}\left[\delta(\eta)+\frac{\tau_0}{8\beta_{\rm T}}\sqrt{\frac{4\beta_{\rm T}-\eta}{\eta}}I_1\left(\frac{\tau_0}{4\beta_{\rm T}}\sqrt{\eta(4\beta_{\rm T}-\eta)}\right)\right]{\rm d}\eta, \\ p_-(\xi){\rm d}\xi &= \frac{\tau_0}{8\beta_{\rm T}}e^{-\tau_0/2}I_0\left(\frac{\tau_0}{4\beta_{\rm T}}\sqrt{(2\beta_{\rm T}-\xi)(\xi+2\beta_{\rm T})}\right){\rm d}\xi,\end{aligned} \tag{A5}$$

where $\eta = (E-E_0)/E_0$ and $\xi = (E_0-E)/E_0$ and E represents the energy of the scattered photon. Rewriting the

probabilities above in terms of the energies gives the expressions (see Eq.(36) in Lyutikov & Gavriil 2006):

$$p_+(E,E_0)\mathrm{d}E_0 = e^{-\tau_0/2}\left[\delta(E-E_0) + \frac{\tau_0}{8\beta_{\rm T}E_0}\sqrt{\frac{4\beta_{\rm T}E_0 - E + E_0}{E-E_0}}\right. \tag{A6}$$

$$\left. I_1\left(\frac{\tau_0}{4\beta_{\rm T}E_0}\sqrt{(E-E_0)(4\beta_{\rm T}E_0 - E + E_0)}\right)\right] dE_0,$$

$$p_-(E,E_0)\mathrm{d}E_0 = \frac{\tau_0}{8\beta_{\rm T}E_0}e^{-\tau_0/2} \tag{A7}$$

$$I_0\left(\frac{\tau_0}{4\beta_{\rm T}E_0}\sqrt{(2\beta_{\rm T}E_0 - E_0 + E)(E_0 - E + 2\beta_{\rm T}E_0)}\right)\mathrm{d}E_0,$$

From the plots in Fig. 2 in Lyutikov & Gavriil (2006) a photon that is reflected can both gain or lose energy. Since the dipolar magnetic field decays as $\sim r^{-3}$, the cyclotron frequency of electrons gyrating along the magnetic field lines also decreases as $\sim r^{-3}$ as one moves far away from the star surface. Therefore if a photon is resonantly scattered and reflected back towards the stellar surface it could interact again with the electron plasma only if it gains energy. If this happens it can be transmitted again towards the observer. Since in the transmission the photon only gain energy, as it propagates far away from the star it will not interact again with electrons since they will have cyclotron frequencies below its energy.

These probabilities depends on two free parameters, i.e., τ_0 and $\beta_{\rm T}$. Following Gullón et al. (2015) we assumed that these two parameters depends on the dipolar magnetic field strength B of the neutron star according to the following scaling relations:

$$\tau_0 = \begin{cases} 0.001 & \text{if } B \leq 10^{13}\,\mathrm{G} \\ \frac{B}{10^{14}\,\mathrm{G}} & \text{if } B > 10^{13}\,\mathrm{G} \end{cases} \tag{A8}$$

$$\beta_{\rm T} = \begin{cases} 0.001 & \text{if } B \leq 10^{13}\,\mathrm{G} \\ 0.3 & \text{if } B > 10^{13}\,\mathrm{G}. \end{cases} \tag{A9}$$

This approximated relations roughly reproduce the correlations between τ_0, $\beta_{\rm T}$ and B found in Rea et al. (2008).

If $n_{\rm BB}(E_0) = I_{\rm BB}(E_0)/E_0$ is the thermal spectrum coming from the neutron star surface in number of photons, where

$$I_{\rm BB}(E) = \frac{2}{h^3c^2}\frac{E^3}{e^{\frac{E}{k_{\rm B}T_\infty}} - 1}, \tag{A10}$$

is a black-body intensity spectrum, the resulting photon intensity spectrum $n_{\rm RCS}(E)$ due to cyclotron resonant scattering can be computed by summing various contributions coming from various reflection events:

$$n_{\rm RCS}(E) = \int dE_0 p_+(E,E_0) n_{\rm BB}(E_0) + \tag{A11}$$

$$+ \int dE_1 p_+(E,E_1) \int dE_0 p_-(E_1,E_0) n_{\rm BB}(E_0) +$$

$$+ \int dE_3 p_+(E,E_3) \int dE_2 p_-(E_3,E_2) \int dE_1 p_-(E_2,E_1)$$

$$\ldots$$

This sum quickly converges and for numerical purposes we will consider 6 reflection terms (see Lyutikov & Gavriil 2006).

B. PREDICTIVE CHECK AND COVERAGE TEST

To evaluate the reliability of the trained posterior estimator in round 5 of Section 5.2 we check its performance on a test dataset of 300 simulations drawn from the proposal restricted prior. In Fig. 11 we show the inference results on

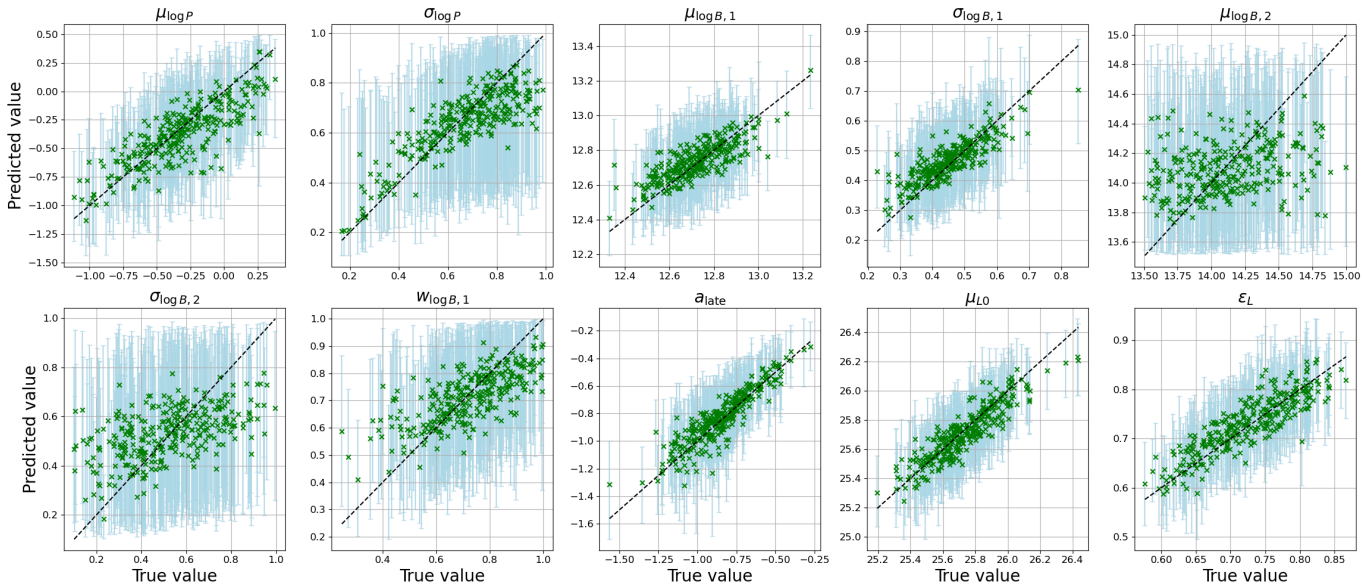

Figure 11. Results of the predictive check test performed on a test dataset of 300 simulations using the density estimator trained on radio pulsars and the sub-populations of young magnetars and XDINSs. We used the density estimator obtained from round 5 (see also Fig. 5). For each parameter, we display the predicted values and their uncertainties, i.e., the median values (green crosses) and the 95% confidence interval (light blue error bars) of the posterior distribution, versus the true value as green crosses and light blue error bars respectively.

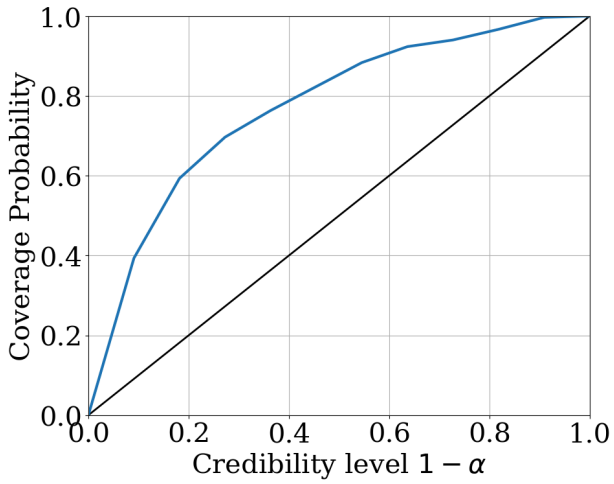

Figure 12. Coverage diagnostic plot for the posterior estimator trained on radio pulsars and the sub-populations of young magnetars and XDINSs from round 5.

the test dataset for each of the 10 parameters we are inferring. For each test sample we report the median value and the 95% credible interval from the predicted posterior distribution and we compare it with the ground truth value. As can be seen by the scatter and uncertainties in the predicted values, the parameters related to the second component of the initial magnetic field, $\mu_{\log B,2}$ and $\sigma_{\log B,2}$ are difficult to constrain. We also performed the coverage probability test (Cook et al. 2006; Hermans et al. 2021) on the trained estimator after the first round of TSNPE. To do this we use the 300 test samples and compute the percentage of ground truth values that fall inside a given credibility level of the corresponding estimated posteriors. The coverage plot shows that the posterior estimator is overall conservative, i.e., the coverage line lies well above the diagonal (see also Graber et al. 2024; Pardo-Araujo et al. 2025, for an explanation of the coverage probability).

REFERENCES

Archibald, R. F., Kaspi, V. M., Tendulkar, S. P., & Scholz, P. 2016, ApJL, 829, L21, doi: 10.3847/2041-8205/829/1/L21

Ascenzi, S., Viganò, D., Dehman, C., et al. 2024, MNRAS, 533, 201, doi: 10.1093/mnras/stae1749

Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, A&A, 558, A33, doi: 10.1051/0004-6361/201322068

Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, AJ, 156, 123, doi: 10.3847/1538-3881/aabc4f

Balucinska-Church, M., & McCammon, D. 1992, ApJ, 400, 699, doi: 10.1086/172032

Barrère, P., Guilet, J., Reboul-Salze, A., Raynaud, R., & Janka, H.-T. 2022, A&A, 668, A79, doi: 10.1051/0004-6361/202244172

Barret, D., & Dupourqué, S. 2024, arXiv e-prints, arXiv:2401.06061, doi: 10.48550/arXiv.2401.06061

Becker, W., & Truemper, J. 1997, A&A, 326, 682, doi: 10.48550/arXiv.astro-ph/9708169

Beloborodov, A. M. 2009, ApJ, 703, 1044, doi: 10.1088/0004-637X/703/1/1044

—. 2013, ApJ, 762, 13, doi: 10.1088/0004-637X/762/1/13

Beniamini, P., Wadiasingh, Z., Hare, J., et al. 2023, MNRAS, 520, 1872, doi: 10.1093/mnras/stad208

Beniamini, P., Wadiasingh, Z., Trigg, A., et al. 2025, ApJ, 980, 211, doi: 10.3847/1538-4357/ada947

Bochenek, C. D., Ravi, V., Belov, K. V., et al. 2020, Nature, 587, 59, doi: 10.1038/s41586-020-2872-x

Borghese, A., & Esposito, P. 2023, in Handbook of X-ray and Gamma-ray Astrophysics, 146, doi: 10.1007/978-981-16-4544-0_102-1

Bradt, H. V., Rothschild, R. E., & Swank, J. H. 1993, A&AS, 97, 355

Caleb, M., Heywood, I., Rajwade, K., et al. 2022, Nature Astronomy, 6, 828, doi: 10.1038/s41550-022-01688-x

Camilo, F., Ransom, S. M., Halpern, J. P., et al. 2006, Nature, 442, 892, doi: 10.1038/nature04986

Canuto, V., Lodenquai, J., & Ruderman, M. 1971, PhRvD, 3, 2303, doi: 10.1103/PhysRevD.3.2303

Carrasco, F., Viganò, D., Palenzuela, C., & Pons, J. A. 2019, MNRAS, 484, L124, doi: 10.1093/mnrasl/slz016

Carreau, T., Fantina, A. F., & Gulminelli, F. 2020, A&A, 640, A77, doi: 10.1051/0004-6361/202038347

Cheng, K. S., Taam, R. E., & Wang, W. 2004, ApJ, 617, 480, doi: 10.1086/425295

Chugunov, A. I., & Horowitz, C. J. 2010, MNRAS, 407, L54, doi: 10.1111/j.1745-3933.2010.00903.x

Cieślar, M., Bulik, T., & Osłowski, S. 2020, MNRAS, 492, 4043, doi: 10.1093/mnras/staa073

Cook, S., Gelman, A., & Rubin, D. 2006, Journal of Computational and Graphical Statistics, 15, doi: 10.1198/106186006X136976

Coti Zelati, F., Rea, N., Pons, J. A., Campana, S., & Esposito, P. 2018, MNRAS, 474, 961, doi: 10.1093/mnras/stx2679

Cranmer, K., Brehmer, J., & Louppe, G. 2020, Proceedings of the National Academy of Science, 117, 30055, doi: 10.1073/pnas.1912789117

Cruise, M., Guainazzi, M., Aird, J., et al. 2025, Nature Astronomy, 9, 36, doi: 10.1038/s41550-024-02416-3

D'Aì, A., Evans, P. A., Burrows, D. N., et al. 2016, MNRAS, 463, 2394, doi: 10.1093/mnras/stw2023

Dask Development Team. 2016, Dask: Library for dynamic task scheduling. http://dask.pydata.org

Dax, M., Green, S. R., Gair, J., et al. 2021, PhRvL, 127, 241103, doi: 10.1103/PhysRevLett.127.241103

De Grandis, D., Taverna, R., Turolla, R., et al. 2021, ApJ, 914, 118, doi: 10.3847/1538-4357/abfdac

Dehman, C., Pons, J. A., Viganò, D., & Rea, N. 2023a, MNRAS, 520, L42, doi: 10.1093/mnrasl/slad003

Dehman, C., Viganò, D., Ascenzi, S., Pons, J. A., & Rea, N. 2023b, MNRAS, 523, 5198, doi: 10.1093/mnras/stad1773

Dehman, C., Viganò, D., Pons, J. A., & Rea, N. 2023c, MNRAS, 518, 1222, doi: 10.1093/mnras/stac2761

Dehman, C., Viganò, D., Rea, N., et al. 2020, ApJL, 902, L32, doi: 10.3847/2041-8213/abbda9

Deistler, M., Goncalves, P. J., & Macke, J. H. 2022, in Advances in Neural Information Processing Systems, ed. S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, & A. Oh, Vol. 35 (Curran Associates, Inc.), 23135–23149. https://proceedings.neurips.cc/paper_files/paper/2022/file/9278abf072b58caf21d48dd670b4c721-Paper-Conference.pdf

Disberg, P., & Mandel, I. 2025, ApJL, 989, L8, doi: 10.3847/2041-8213/adf286

Doroshenko, V. 2024, arXiv e-prints, arXiv:2403.03127, doi: 10.48550/arXiv.2403.03127

Duncan, R. C., & Thompson, C. 1992, ApJL, 392, L9, doi: 10.1086/186413

Edwards, R. T., Bailes, M., Van Straten, W., & Britton, M. C. 2001, Monthly Notices of the Royal Astronomical Society, 326, 358

Esposito, P., Rea, N., & Israel, G. L. 2021, in Astrophysics and Space Science Library, Vol. 461, Astrophysics and Space Science Library, ed. T. M. Belloni, M. Méndez, & C. Zhang, 97–142, doi: 10.1007/978-3-662-62110-3_3

Faucher-Giguère, C.-A., & Kaspi, V. M. 2006, ApJ, 643, 332, doi: 10.1086/501516
Gardenier, D. W., & van Leeuwen, J. 2021, A&A, 651, A63, doi: 10.1051/0004-6361/202040119
Gavriil, F. P., Gonzalez, M. E., Gotthelf, E. V., et al. 2008, Science, 319, 1802, doi: 10.1126/science.1153465
Gehrels, N., Chincarini, G., Giommi, P., et al. 2004, ApJ, 611, 1005, doi: 10.1086/422091
Gençali, A. A., & Ertan, Ü. 2024, MNRAS, 534, 1481, doi: 10.1093/mnras/stae2177
Glorot, X., & Bengio, Y. 2010, in Proceedings of Machine Learning Research, Vol. 9, Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, ed. Y. W. Teh & M. Titterington (Chia Laguna Resort, Sardinia, Italy: PMLR), 249–256. https://proceedings.mlr.press/v9/glorot10a.html
Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759, doi: 10.1086/427976
Gotthelf, E. V. 2003, ApJ, 591, 361, doi: 10.1086/375124
Graber, V., Ronchi, M., Pardo-Araujo, C., & Rea, N. 2024, ApJ, 968, 16, doi: 10.3847/1538-4357/ad3e78
Greenberg, D. S., Nonnenmacher, M., & Macke, J. H. 2019, arXiv e-prints, arXiv:1905.07488, doi: 10.48550/arXiv.1905.07488
Gullón, M., Miralles, J. A., Viganò, D., & Pons, J. A. 2014, MNRAS, 443, 1891, doi: 10.1093/mnras/stu1253
Gullón, M., Pons, J. A., Miralles, J. A., et al. 2015, MNRAS, 454, 615, doi: 10.1093/mnras/stv1644
Harding, A. K. 2013, Frontiers of Physics, 8, 679, doi: 10.1007/s11467-013-0285-0
Harris, C. R., Jarrod Millman, K., van der Walt, S. J., et al. 2020, Nature, 585, 357, doi: 10.1038/s41586-020-2649-2
He, K., Zhang, X., Ren, S., & Sun, J. 2015, preprint, arXiv:1502.01852
Hermans, J., Delaunoy, A., Rozet, F., et al. 2021, arXiv e-prints, arXiv:2110.06581, doi: 10.48550/arXiv.2110.06581
Hunter, J. D. 2007, Comput. Sci. Eng., 9, 90, doi: 10.1109/MCSE.2007.55
Inserra, C., Smartt, S. J., Jerkstrand, A., et al. 2013, ApJ, 770, 128, doi: 10.1088/0004-637X/770/2/128
Jacoby, B. A., Bailes, M., Ord, S. M., Edwards, R. T., & Kulkarni, S. R. 2009, ApJ, 699, 2009, doi: 10.1088/0004-637X/699/2/2009
Johnston, S., Smith, D. A., Karastergiou, A., & Kramer, M. 2020, MNRAS, 497, 1957, doi: 10.1093/mnras/staa2110
Jones, E., Oliphant, T. E., Peterson, P., et al. 2001, SciPy: Open source scientific tools for Python. http://www.scipy.org/
Kaspi, V. M. 2010, Proceedings of the National Academy of Science, 107, 7147, doi: 10.1073/pnas.1000812107
Kaspi, V. M., & Beloborodov, A. M. 2017, ARA&A, 55, 261, doi: 10.1146/annurev-astro-081915-023329
Keane, E., Bhattacharyya, B., Kramer, M., et al. 2015, in Advancing Astrophysics with the Square Kilometre Array (AASKA14), 40, doi: 10.22323/1.215.0040
Keane, E. F., & Kramer, M. 2008, MNRAS, 391, 2009, doi: 10.1111/j.1365-2966.2008.14045.x
Keane, E. F., Graber, V., Levin, L., et al. 2025, The Open Journal of Astrophysics, 8, 54256, doi: 10.33232/001c.154256
Keith, M. J., Jameson, A., van Straten, W., et al. 2010, MNRAS, 409, 619, doi: 10.1111/j.1365-2966.2010.17325.x
Kingma, D. P., & Ba, J. 2014, arXiv e-prints, arXiv:1412.6980, doi: 10.48550/arXiv.1412.6980
Kurpas, J., Pires, A. M., Schwope, A. D., et al. 2026, A&A, 705, A148, doi: 10.1051/0004-6361/202556815
Lam, S. K., Pitrou, A., & Seibert, S. 2015, in Proc. Second Workshop on the LLVM Compiler Infrastructure in HPC, 1–6, doi: 10.1145/2833157.2833162
Lander, S. K. 2023, ApJL, 947, L16, doi: 10.3847/2041-8213/acca1f
Li, C., Zhao, G., Jia, Y., et al. 2019, ApJ, 871, 208, doi: 10.3847/1538-4357/aafa17
Lorimer, D. R., & Kramer, M. 2012, Handbook of Pulsar Astronomy (Cambridge University Press)
Lorimer, D. R., Faulkner, A. J., Lyne, A. G., et al. 2006, MNRAS, 372, 777, doi: 10.1111/j.1365-2966.2006.10887.x
Lü, H.-J., & Zhang, B. 2014, ApJ, 785, 74, doi: 10.1088/0004-637X/785/1/74
Lueckmann, J.-M., Goncalves, P. J., Bassetto, G., et al. 2017, arXiv e-prints, arXiv:1711.01861, doi: 10.48550/arXiv.1711.01861
Lyutikov, M., & Gavriil, F. P. 2006, MNRAS, 368, 690, doi: 10.1111/j.1365-2966.2006.10140.x
Makarenko, E. I., Igoshev, A. P., & Kholtygin, A. F. 2021, MNRAS, 504, 5813, doi: 10.1093/mnras/stab1175
Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, AJ, 129, 1993, doi: 10.1086/428488
Manchester, R. N., Lyne, A. G., Camilo, F., et al. 2001, MNRAS, 328, 17, doi: 10.1046/j.1365-8711.2001.04751.x
Marino, A., Dehman, C., Kovlakas, K., et al. 2024, Nature Astronomy, 8, 1020, doi: 10.1038/s41550-024-02291-y
McKinney, W. 2010, in Proceedings of the 9th Python in Science Conference, ed. S. van der Walt & J. Millman, 51–56. http://conference.scipy.org/proceedings/scipy2010/mckinney.html

Metzger, B. D., Giannios, D., Thompson, T. A., Bucciantini, N., & Quataert, E. 2011, MNRAS, 413, 2031, doi: 10.1111/j.1365-2966.2011.18280.x
Mishra-Sharma, S., & Cranmer, K. 2022, PhRvD, 105, 063017, doi: 10.1103/PhysRevD.105.063017
Morello, V., Barr, E. D., Stappers, B. W., Keane, E. F., & Lyne, A. G. 2020, MNRAS, 497, 4654, doi: 10.1093/mnras/staa2291
Narayan, R., & Ostriker, J. P. 1990, ApJ, 352, 222, doi: 10.1086/168529
Nobili, L., Turolla, R., & Zane, S. 2008, MNRAS, 386, 1527, doi: 10.1111/j.1365-2966.2008.13125.x
Oliphant, T. E. 2006, A guide to NumPy (USA: Trelgol Publishing)
Page, D., Lattimer, J. M., Prakash, M., & Steiner, A. W. 2004, ApJS, 155, 623, doi: 10.1086/424844
Papamakarios, G., & Murray, I. 2016, arXiv e-prints, arXiv:1605.06376, doi: 10.48550/arXiv.1605.06376
Pardo-Araujo, C., Rea, N., Ronchi, M., & Graber, V. 2026, arXiv e-prints, arXiv:2601.16159, doi: 10.48550/arXiv.2601.16159
Pardo-Araujo, C., Ronch, M., Graber, V., & Rea, N. 2025, A&A, 696, A114, doi: 10.1051/0004-6361/202453314
Pastor-Marazuela, I., van Leeuwen, J., Bilous, A., et al. 2025, A&A, 693, A279, doi: 10.1051/0004-6361/202450953
Paszke, A., Gross, S., Massa, F., et al. 2019, preprint, arXiv:1912.01703. https://arxiv.org/abs/1912.01703
Pearson, J. M., Chamel, N., Potekhin, A. Y., et al. 2018, MNRAS, 481, 2994, doi: 10.1093/mnras/sty2413
Perez, F., & Granger, B. E. 2007, Comput. Sci. Eng., 9, 21, doi: 10.1109/MCSE.2007.53
Perna, R., & Pons, J. A. 2011, ApJL, 727, L51, doi: 10.1088/2041-8205/727/2/L51
Pétri, J. 2019, MNRAS, 485, 4573, doi: 10.1093/mnras/stz711
Philippov, A., Tchekhovskoy, A., & Li, J. G. 2014, MNRAS, 441, 1879, doi: 10.1093/mnras/stu591
Pons, J. A., Viganò, D., & Rea, N. 2013, Nature Physics, 9, 431, doi: 10.1038/nphys2640
Popov, S. B. 2023, Universe, 9, 273, doi: 10.3390/universe9060273
Popov, S. B., Pons, J. A., Miralles, J. A., Boldin, P. A., & Posselt, B. 2010, MNRAS, 401, 2675, doi: 10.1111/j.1365-2966.2009.15850.x
Posselt, B., Karastergiou, A., Johnston, S., et al. 2023, MNRAS, 520, 4582, doi: 10.1093/mnras/stac3383
Potekhin, A. Y., Pons, J. A., & Page, D. 2015, SSRv, 191, 239, doi: 10.1007/s11214-015-0180-9
Qu, Y., & Bransgrove, A. 2026, ApJ, 998, 190, doi: 10.3847/1538-4357/ae3a9d
Rea, N., Borghese, A., Esposito, P., et al. 2016, ApJL, 828, L13, doi: 10.3847/2041-8205/828/1/L13
Rea, N., & De Grandis, D. 2025, arXiv e-prints, arXiv:2503.04442, doi: 10.48550/arXiv.2503.04442
Rea, N., & Esposito, P. 2011, in Astrophysics and Space Science Proceedings, Vol. 21, High-Energy Emission from Pulsars and their Systems, ed. D. F. Torres & N. Rea, 247, doi: 10.1007/978-3-642-17251-9_21
Rea, N., Gullón, M., Pons, J. A., et al. 2015, ApJ, 813, 92, doi: 10.1088/0004-637X/813/2/92
Rea, N., Hurley-Walker, N., & Caleb, M. 2026, arXiv e-prints, arXiv:2601.10393, doi: 10.48550/arXiv.2601.10393
Rea, N., Zane, S., Turolla, R., Lyutikov, M., & Götz, D. 2008, ApJ, 686, 1245, doi: 10.1086/591264
Ronchi, M., Graber, V., Garcia-Garcia, A., Rea, N., & Pons, J. A. 2021, ApJ, 916, 100, doi: 10.3847/1538-4357/ac05bd
Ronchi, M., Rea, N., Graber, V., & Hurley-Walker, N. 2022, ApJ, 934, 184, doi: 10.3847/1538-4357/ac7cec
Rozwadowska, K., Vissani, F., & Cappellaro, E. 2021, NewA, 83, 101498, doi: 10.1016/j.newast.2020.101498
Santangelo, A., Zane, S., Feng, H., et al. 2019, Science China Physics, Mechanics, and Astronomy, 62, 29505, doi: 10.1007/s11433-018-9234-3
Sautron, M., McEwen, A. E., Younes, G., et al. 2025, ApJ, 986, 88, doi: 10.3847/1538-4357/add0aa
Sengar, R., Bailes, M., Balakrishnan, V., et al. 2025, MNRAS, 536, 3159, doi: 10.1093/mnras/stae2716
Spitkovsky, A. 2006, ApJL, 648, L51, doi: 10.1086/507518
Stratta, G., Dainotti, M. G., Dall'Osso, S., Hernandez, X., & De Cesare, G. 2018, ApJ, 869, 155, doi: 10.3847/1538-4357/aadd8f
Tejero-Cantero, A., Boelts, J., Deistler, M., et al. 2020, The Journal of Open Source Software, 5, 2505, doi: 10.21105/joss.02505
Tong, H., Xu, R. X., Song, L. M., & Qiao, G. J. 2013, ApJ, 768, 144, doi: 10.1088/0004-637X/768/2/144
Truemper, J. 1982, Advances in Space Research, 2, 241, doi: 10.1016/0273-1177(82)90070-9
Turolla, R. 2009, in Astrophysics and Space Science Library, Vol. 357, Astrophysics and Space Science Library, ed. W. Becker, 141, doi: 10.1007/978-3-540-76965-1_7
Turolla, R., Zane, S., & Watts, A. L. 2015, Reports on Progress in Physics, 78, 116901, doi: 10.1088/0034-4885/78/11/116901
van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, Comput. Sci. Eng., 13, 22, doi: 10.1109/MCSE.2011.37

van Leeuwen, J., Kooistra, E., Oostrum, L., et al. 2023, A&A, 672, A117, doi: 10.1051/0004-6361/202244107
Vasist, M., Rozet, F., Absil, O., et al. 2023, A&A, 672, A147, doi: 10.1051/0004-6361/202245263
Verberne, S., & Vink, J. 2021, MNRAS, 504, 1536, doi: 10.1093/mnras/stab940
Viganò, D., Garcia-Garcia, A., Pons, J. A., Dehman, C., & Graber, V. 2021, Computer Physics Communications, 265, 108001, doi: 10.1016/j.cpc.2021.108001
Viganò, D., & Pons, J. A. 2012, MNRAS, 425, 2487, doi: 10.1111/j.1365-2966.2012.21679.x
Viganò, D., Rea, N., Pons, J. A., et al. 2013, MNRAS, 434, 123, doi: 10.1093/mnras/stt1008
Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, Nature Methods, 17, 261, doi: 10.1038/s41592-019-0686-2
Voges, W., Aschenbach, B., Boller, T., et al. 1999, A&A, 349, 389, doi: 10.48550/arXiv.astro-ph/9909315
Wadiasingh, Z., & Timokhin, A. 2019, ApJ, 879, 4, doi: 10.3847/1538-4357/ab2240
Wainscoat, R. J., Cohen, M., Volk, K., Walker, H. J., & Schwartz, D. E. 1992, ApJS, 83, 111, doi: 10.1086/191733
Wang, Y., & van Leeuwen, J. 2024, A&A, 690, A377, doi: 10.1051/0004-6361/202450673
Watson, M. G., Auguères, J.-L., Ballet, J., et al. 2001, A&A, 365, L51, doi: 10.1051/0004-6361:20000067
Weisskopf, M. C., Brinkman, B., Canizares, C., et al. 2002, PASP, 114, 1, doi: 10.1086/338108
Willingale, R., Starling, R. L. C., Beardmore, A. P., Tanvir, N. R., & O'Brien, P. T. 2013, MNRAS, 431, 394, doi: 10.1093/mnras/stt175
Wilms, J., Allen, A., & McCray, R. 2000, ApJ, 542, 914, doi: 10.1086/317016
Xu, Y.-J., Peng, H.-L., Weng, S.-S., Zhang, X., & Ge, M.-Y. 2025, ApJ, 981, 100, doi: 10.3847/1538-4357/adaebc
Yamasaki, S., Lyubarsky, Y., Granot, J., & Göğüş, E. 2020, MNRAS, 498, 484, doi: 10.1093/mnras/staa2223
Yao, J. M., Manchester, R. N., & Wang, N. 2017, ApJ, 835, 29, doi: 10.3847/1538-4357/835/1/29
Zammit-Mangion, A., Sainsbury-Dale, M., & Huser, R. 2025, Annual Review of Statistics and Its Application, 12, 311, doi: 10.1146/annurev-statistics-112723-034123
Zane, S., Rea, N., Turolla, R., & Nobili, L. 2009, MNRAS, 398, 1403, doi: 10.1111/j.1365-2966.2009.15190.x
Zhang, B., & Hu, R.-C. 2025, ApJL, 994, L20, doi: 10.3847/2041-8213/ae1023
Zhang, L., Ridolfi, A., Li, D., et al. 2025, arXiv e-prints, arXiv:2512.17214, doi: 10.48550/arXiv.2512.17214
Zhou, X., Huang, H.-T., Cheng, Q., & Zheng, X.-P. 2024, MNRAS, 530, 1636, doi: 10.1093/mnras/stae954
Zonca, A., Singer, L., Lenz, D., et al. 2019, The Journal of Open Source Software, 4, 1298, doi: 10.21105/joss.01298