UAB
Universitat Autònoma de Barcelona

UNIVERSITAT AUTÒNOMA DE BARCELONA
DEPARTAMENT DE FÍSICA

U🔺B
Universitat Autònoma
de Barcelona

DOCTORAL THESIS

Population synthesis of Galactic pulsars with machine learning

Author:
Michele Ronchi

Advisors:
Prof. Nanda Rea
Dr. Vanessa Graber

Tutor:
Prof. Lluís Font Guiteras

A thesis submitted for the degree of
Doctor of Philosophy in Physics

Institute of Space Sciences (ICE-CSIC/IEEC)

Barcelona, December 2023

Abstract

The recent advancements in data availability and computational power have allowed the development of new machine-learning algorithms and data-analysis techniques that have opened up new avenues to study the Galactic neutron-star population.

This thesis work represents the first efforts to combine population synthesis studies of the Galactic isolated neutron stars with deep-learning techniques with the aim of better understanding neutron-star birth properties and evolution.

In particular, we develop a flexible population-synthesis framework to model the dynamical and magneto-rotational evolution of neutron stars, their emission in radio and their detection with radio telescopes. We first study the feasibility of using deep neural networks to infer the dynamical properties at birth of a simulated population of neutron stars from density maps storing the information on their sky position and proper motion. We then explore a simulation-based inference approach to constrain our physical models from the observed radio pulsar population. We employ a neural density estimator to predict the posterior distribution of the parameters describing the birth magnetic-field and spin-period distributions and the late-time magnetic-field decay from density maps containing information on the spin period and spin-period derivative. Our results for the initial magneto-rotational properties agree with the findings of previous works while we constrain the late-time evolution of the magnetic field in neutron stars for the first time.

Besides the development of the population-synthesis framework, this thesis also studies possible scenarios to explain the puzzling nature of recently discovered periodic radio sources with very long periods of the order of thousands of seconds.

In particular, by assuming a neutron-star origin, we study the spin-period evolution of a newborn neutron star interacting with a supernova fallback disk and find that the combination of strong, magnetar-like magnetic fields and moderate accretion rates can lead to very large spin periods on timescales of ten thousands of years. Moreover, we perform population synthesis studies to assess the possibility for these sources to be either neutron stars or magnetic white dwarfs emitting coherently through magnetic dipolar losses. These discoveries have opened up a new perspective on the neutron-star population and have started to question our current understanding of how coherent radio emission is produced in pulsar magnetospheres.

Overall this thesis work represents the first step towards developing a simulation framework able to model the different neutron star classes in a unified scenario and constraining the properties of the neutron-star population as a whole from multi-wavelength observations through machine learning.

iv

This thesis is based on the following publications:

- Ronchi, Graber, Garcia-Garcia, Rea, and Pons (2021)
  "Analyzing the Galactic Pulsar Distribution with Machine Learning"
  Published in: ApJ 916.2, 100, p. 100

- Ronchi, Rea, Graber, and Hurley-Walker (2022)
  "Long-period Pulsars as Possible Outcomes of Supernova Fallback Accretion"
  Published in: ApJ 934.2, 184, p. 184

- Rea, Coti Zelati, Dehman, Hurley-Walker, de Martino, Bahramian, Buckley, Brink, Kawka, Pons, Viganò, Graber, Ronchi, Pardo Araujo, Borghese, Parent, and Galvin (2022)
  "Constraining the Nature of the 18 min Periodic Radio Transient GLEAM-X J162759.5-523504.3 via Multiwavelength Observations and Magneto-thermal Simulations"
  Published in: ApJ 940.1, 72, p. 72

- Rea, Hurley-Walker, Pardo-Araujo, Ronchi, Graber, Coti Zelati, De Martino, Bahramian, McSweeney, Galvin, Hyman, and Dall'Ora (2023)
  "A long-period radio transient active for three decades: population study in the neutron star and white dwarf rotating dipole scenarios"
  Accepted for publication in ApJ

- Graber, Ronchi, Pardo-Araujo and Rea (2023)
  "Simulation-based inference for pulsar population synthesis"
  Submitted to ApJ

Acknowledgements

These past four years, besides being challenging, have been very enjoyable, stimulating and rewarding in a lot of aspects and I am very grateful to all the people that I met and who shared this part of my life with me.

First of all I would like to deeply thank my supervisors Nanda and Vanessa. Your guidance, exceptional expertise, enthusiasm, support and understanding have been fundamental during these years. I learnt a lot from you both from a scientific and a human point of view. Thanks to you I became more aware of both the positive aspects and drawbacks of academia and your advices are going to be incredibly valuable for choosing my future path.

People who have previously been and are currently part of the MAGNESIA group are all fantastic. The atmosphere has always been very friendly and supportive. Thanks to the entire Italian crew (Alice, Alessio, Francesco, Stefano, Daniele, Arianna and the new-comers Davide, Thomas and Anna) who always made me feel like at home and to all the people coming from Spain and other parts of the world (Albert I and Albert II, Celsa, Clara, Abu, Emilie, Kostas, Christine and Raj) who opened up my horizons. I feel enriched of having met so many interesting people with their diverse perspectives, vast knowledge and experience. I learned a lot from all of you. Special thanks to my PhD mates, Abu for your kindness and wisdom in every situation, Clara for your incredible determination and Celsa, also team-mate (and Spanish teacher), for your contagious enthusiasm and authenticity in everything you do. Together with Vanessa, the harmony and synergy in our population-synthesis group was fantastic and this thesis is only possible thanks to such an amazing team-work.

I extend my gratitude to all the people I met at ICE, a very international and stimulating environment with a lot of PhD students sharing the same journey, and people at PIC, for all the helpful conversations on the computing side. I would further like to thank my examiners, Ben Stappers, Alessandro Patruno and Helena Dominguez Sanchez, who took their time to read this thesis and give constructive feedback that helped me to improve this work.

Lastly, I would like to thank my family, my mum, dad, brother and my grandmother and all my friends back home for their support, encouragement and understanding in all these years.

Contents

List of Figures

List of Tables

Physical and Astronomical Constants

Speed of light in a vacuum	$c = 2.99792458 \times 10^{10} \, \mathrm{cm \, s^{-1}}$
Gravitational constant	$G = 6.67259 \times 10^{-8} \, \mathrm{cm^3 \, g^{-1} \, s^{-2}}$
Planck constant	$h = 6.6260755 \times 10^{-27} \, \mathrm{erg \, s}$
Electron charge	$e = 4.8032068 \times 10^{-10} \, \mathrm{esu}$
Electron mass	$m_e = 9.1093897 \times 10^{-28} \, \mathrm{g}$
Proton mass	$m_p = 1.6726231 \times 10^{-24} \, \mathrm{g}$
Neutron mass	$m_n = 1.6749286 \times 10^{-24} \, \mathrm{g}$
Boltzmann constant	$k_{\mathrm{B}} = 1.380658 \times 10^{-16} \, \mathrm{erg \, K^{-1}}$
Thomson cross-section	$\sigma_{\mathrm{T}} = 6.65245 \times 10^{-25} \, \mathrm{cm^{-2}}$
Solar mass	$M_\odot = 1.989 \times 10^{33} \, \mathrm{g}$
Parsec	$\mathrm{pc} = 3.0856 \times 10^{18} \, \mathrm{cm}$
Jansky	$\mathrm{Jy} = 1.0 \times 10^{-23} \, \mathrm{erg \, cm^{-2} \, s^{-1} \, Hz^{-1}}$
Year	$\mathrm{yr} = 3.1557 \times 10^{7} \, \mathrm{s}$

Chapter 1

Neutron stars

1.1 Introduction

The idea of the possible existence of neutron stars was first proposed in 1934, two years after the discovery of the neutron itself (Chadwick, 1932), when Baade and Zwicky (1934) wondered if a celestial body made entirely of neutrons might remain after a supernova explosion. It was not until the 1960s that Baade and Zwicky's hypothetical neutron stars were finally detected. While working as a PhD student at the University of Cambridge, Jocelyn Bell Burnell noticed a strange, regularly pulsed radio signal in the data collected using the Mullard Radio Astronomy Observatory. These signals were initially interpreted as the result of radial pulsations in either a white dwarf or a neutron star (Hewish et al., 1968). This groundbreaking discovery opened up new avenues for exploring the intriguing relationship between neutron stars and their potential to emit highly focused beams of electromagnetic radiation. Soon after it was suggested by Pacini (1968) and Gold (1968) that coherent emission could occur in rotating neutron stars with strong magnetic fields where particles could be accelerated up to relativistic speeds and emit beams of electromagnetic radiation. This can be detected as pulsed signals as the neutron star rotates like a lighthouse. Hence the name "pulsar" was coined. By that time about 20 similar objects had been identified including a source in the Crab nebula (Staelin and Reifenstein, 1968; Comella et al., 1969) and one in the Vela supernova remnant (Large, Vaughan, and Mills, 1968). The identification of these pulsars with rotating neutron stars was in the end confirmed by the measurement of a spin-down in the period of the Crab pulsar (Richards and Comella, 1969) as a spin-down is more easily achieved in rotating objects rather than in pulsating ones. Furthermore, the identification of some pulsars with supernova remnants further supported the initial idea of Baade and Zwicky.

With subsequent technological advancements in observational astronomy, the population of known neutron stars has grown exponentially. Since the first detection, around 3000 pulsars have been discovered, not only in the radio band (Manchester et al., 2005)[1]. Pulsars can emit across a wider wavelength range from radio to gamma-rays and even be radio-quiet. However, an enduring challenge lies in accurately characterising their birth properties as well as deciphering the mechanisms that regulate their formation and evolution. Untangling the contributing factors, such as the initial magnetic field strength and spin period distributions, holds paramount importance in understanding the diverse outcomes of stellar evolution and the underlying physical processes within core-collapse supernovae. By combining observational constraints with data analysis and sophisticated numerical simulations, researchers can refine existing theoretical models and shed light on the

---

[1]see the ATNF pulsar catalog https://www.atnf.csiro.au/research/pulsar/psrcat/

complex physical processes involved in the formation and evolution of these dense remnants.

In this chapter we will explore the basics of neutron star physics and our current understanding of how they evolve, emit electromagnetic radiation, with particular focus on the radio emission, and how we can detect them (Sections from 1.2 to 1.12). Additionally, we will provide a summary of the challenges involved in modelling the Galactic population of neutron stars (Sections 1.13 and 1.14) and explore how employing a population synthesis framework can help in determining their overall properties (Section 1.15).

1.2 Formation

The progenitors of most neutron stars are massive stars of spectral types O and B, characterised by initial masses between around 8 and 25 $M_\odot$ (Shapiro and Teukolsky, 1983). In general, depending on its initial mass, a star spend from a few million years up to more than ten billion years on the main sequence, burning hydrogen into helium (Phillips, 1994; Prialnik, 2000). In this stage, the thermal pressure released by the thermonuclear fusion reactions in its core supports the star against gravitational collapse. However, the more massive a star is, the faster it consumes its nuclear fuel and the faster it evolves towards its death. After exhausting all the hydrogen in the core the thermal pressure support coming from hydrogen burning vanishes. The inner regions of the star, now mainly composed of helium "ashes", start to contract due to their own gravity. As a consequence they heat up while the outer envelope expands and cools down. The star becomes a red giant. As the temperature in the helium core grows upon contraction, new nuclear fusion reactions are triggered. Helium starts to burn and to be converted into carbon, oxygen and neon. The released energy allows to restore again the equilibrium with gravity and the star enters a new stable phase that usually lasts one-tenth of the main sequence phase. However, once the helium fuel is exhausted, the core (now composed mainly of oxygen and carbon) starts to contract again.

In stars with an initial mass of less than about 8 $M_\odot$, the dense core reaches equilibrium in a new state of matter called a degenerate electron gas. In this state, the stellar core is supported against further gravitational contraction thanks to the pressure exerted by the electrons due to the Pauli exclusion principle, even in the absence of nuclear reactions. In the meanwhile, the outer envelopes of the star expand further and are gradually dispersed. What is left is the exposed core of the star with a mass of $\sim 0.5 - 1\,M_\odot$ and a radius of $\sim 10^8$ cm, made primarily by helium, carbon and oxygen. A white dwarf has formed.

For stars with an initial mass greater than about 8 $M_\odot$ the electron degeneracy pressure is not enough to counteract gravity. Therefore they continue the sequence of core contraction and ignition of new nuclear reactions. As time goes by the inner core reaches temperatures and densities favourable for the synthesis of progressively heavier elements. After hydrogen and helium, these massive stars burn carbon, neon, oxygen, and silicon on progressively faster timescales. In this way, they undergo all the stages of nuclear burning up to the production of elements in the "iron group" with atomic number $A = 56$. The iron group elements are the most tightly bound nuclei and their synthesis into heavier elements consumes, rather than releases, energy. Since no further energy can be released from the nuclear fusion reactions as the iron core grows it becomes unstable and is not further supported against gravitational collapse. As the iron core approaches the Chandrasekhar mass limit

of about 1.4 $M_\odot$ (Chandrasekhar, 1931), the central temperature and density reaches values of about $\sim 10^9$ K and $\sim 10^9$ g cm^{-3} respectively. The temperature is high enough that nuclear photo-disintegration is triggered: the abundant energetic photons trapped in the hot dense core are partly consumed to unbind nuclei into lighter elements in endothermic reactions. In such a way the photon radiation pressure is lost. Furthermore, at such densities, reactions like electron capture $e^- + p \longrightarrow n + \nu_e$ become energetically favourable. The electrons are absorbed by the nuclei which become very neutron-rich in what is called neutronisation. This process depletes the core of electrons and their supporting degeneracy pressure and produces many escaping neutrinos which carry away a big amount of energy. This leads to an almost unrestrained collapse of the stellar core on a free-fall timescale of a few seconds.

As the collapse proceeds, the core central density increases to the point that neutrons start to drip out of the nuclei and form a degenerate gas. At densities $\gtrsim 10^{14}$ g cm^{-3}, the nuclei are completely dissolved into homogeneous nuclear matter mostly composed of neutrons. If the mass of the collapsing core is not too high ($\lesssim 3\,M_\odot$), the degeneracy pressure exerted by the neutrons is finally able to halt the collapse. This triggers a shock wave, that propagates outwards through the collapsing stellar layers and blows them up in what is observed as a supernova explosion. What is left is a very hot ($T \sim 10^8$ K) and compact remnant of mass $M \sim 1 - 2$ $M_\odot$ and radius $R \sim 10^6$ cm: a neutron star.

What described above is the evolution of an isolated massive star. Another channel that leads to the formation of a neutron star is through accretion induced collapse. In this scenario a white dwarf in a binary system can collapse upon surpassing the Chandrasekhar mass limit due to mass accretion from a companion star (Fryer et al., 1999).

1.3 Internal structure and the mass-radius relation

To our knowledge, neutron stars are the most compact objects in the Universe where matter can still reach a stable configuration under the effect of extreme gravity. Their compactness, defined by the ratio between their Schwarzschild radius $R_{\mathrm{S}} = \frac{2GM_{\mathrm{NS}}}{c^2}$ and their radius R_{NS}, is of the order:

$$\frac{R_{\mathrm{S}}}{R_{\mathrm{NS}}} \simeq 0.376 \left(\frac{M_{\mathrm{NS}}}{1.4 M_\odot} \right) \left(\frac{R_{\mathrm{NS}}}{11\,\mathrm{km}} \right)^{-1}. \tag{1.1}$$

This is close to the most extreme value given by a Schwarzschild black hole where the compactness is 1. The density grows from values less than 10^4 g cm^{-3} in the outermost envelope to nuclear density values $10^{14} - 10^{15}$ g cm^{-3} in the neutron star core. As you move radially into the star from the surface, the pressure increases to support the increasing weight of the material above any given point. Hence the structure and state of matter change as we move inside (see Chamel and Haensel, 2008, for a review). The outer layers (~ 100 m) are a liquid ocean of heavy nuclei of the iron group and degenerate free electrons. As one moves into the crust (~ 1 km), the nuclei are more packed together and the Coulomb interactions between them begin to dominate over the thermal energy. The nuclei rearrange themselves to form a solid elastic lattice. On the other end, the degenerate free electrons become relativistic. At the basis of the crust and the interface with the core, the heavy nuclei become more and more neutron-rich and deformed into cylindrical and planar shapes. This is the so-called "nuclear pasta phase". Moving into the inner core the nuclei are completely dissolved into an ocean of free neutrons togeteher with protons, ultra-relativistic
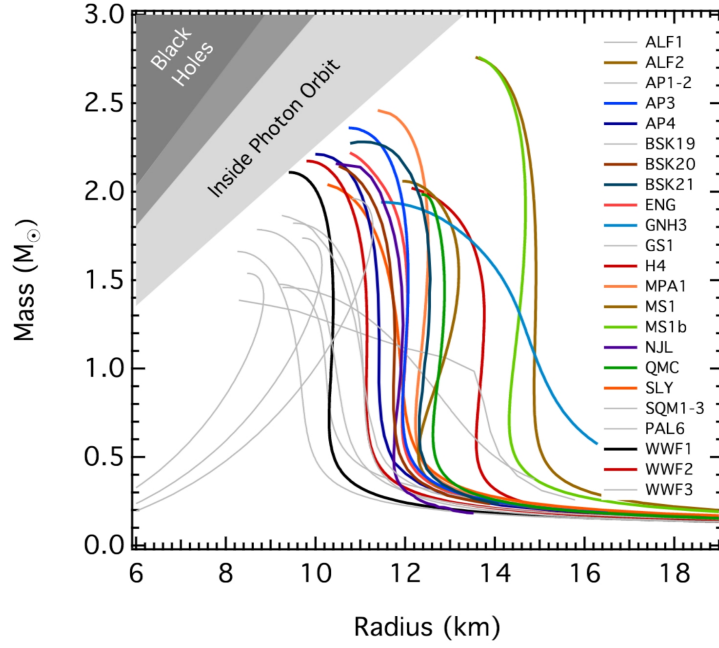
FIGURE 1.1: Neutron star mass-radius curves corresponding to different equations of state calculated under different physical assumptions and using a range of computational approaches (see Özel and Freire, 2016, for more details) (Figure taken from Özel and Freire, 2016).

electrons and muons. The neutron degenerate pressure supports the neutron star against collapse. However, at these extreme densities, the exact composition and the properties of matter in the inner core are not clear as exotic particles may be present. Therefore the exact equation of state, i.e. the relation linking pressure and density, for neutron stars is still unknown. Oppenheimer and Volkoff (1939) and Tolman (1939) were the first ones to propose an equation of state of a relativistic gas of degenerate neutrons and a model for the hydrostatic equilibrium of general relativistic bodies also known as Oppenheimer-Volkoff-Tolman equation:

$$\frac{\mathrm{d}p(r)}{\mathrm{d}r} = \frac{GM(r)\rho(r)}{r^2} \left[1 + \frac{p(r)}{\rho(r)c^2} \right] \left[1 + \frac{4\pi r^3 p(r)}{M(r)c^2} \right] \left[1 - \frac{2GM(r)}{c^2 r} \right]^{-1}, \qquad (1.2)$$

where $p(r)$ and $\rho(r)$ are the pressure and density, respectively, and $M(r)$ is the mass within the radius r. In particular, the solution of this equation maps a given equation of state to the corresponding mass-radius relation, linking in this way the microscopic to the macroscopic properties of the star. Since then the knowledge of the behaviour of matter at nuclear densities has progressed and several models for the equation of state of neutron stars have been proposed. The uncertainty is still big. For a given mass, these equations of state allow a range of radii between $\sim 7\,\mathrm{km}$ up to $\sim 16\,\mathrm{km}$ (see Figure 1.1). Therefore, to constrain the equation of state for the neutron star matter, it is crucial to obtain independent measurements of the mass and radius (see Özel and Freire, 2016, for a review).

The currently available measurements of neutron star masses come primarily from the study of neutron stars in binary systems. Up to now the minimum mass measured for a neutron star is $1.174 \pm 0.004\,M_\odot$ for PSR J0453+1559 (Martinez et al., 2015) while the heaviest neutron star is PSR J0952–0607 which has an estimated

mass of $2.35 \pm 0.17\, M_\odot$ (Romani et al., 2022). Finding the maximum mass for neutron stars is of particular relevance since it can rule out the equations of state that predict maximum masses below this value.

Until recently, the measurement of neutron-star radii mainly relied on X-ray observations and spectral analysis of quiescent and accreting neutron stars (Miller et al., 2019; Riley et al., 2019). The typical estimates for the radii obtained with these methods are in the range $10 - 14$ km (see, e.g., Steiner et al., 2018). However, these analyses are often affected by systematics that are difficult to quantify. After the serendipitous detection of the famous binary neutron-star merger GW170817 and its electromagnetic counterparts (Abbott et al., 2017b; Abbott et al., 2017c; Abbott et al., 2017a), the beginning of the multi-messenger observation era has opened a new window on this front. For example, by combining electromagnetic and gravitational wave data Margalit and Metzger (2017) managed to put constraints on the maximum neutron star mass. Capano et al. (2020) and Raaijmakers et al. (2021) combined nuclear theory with the data from the gravitational waves and electromagnetic emission from this event using a Bayesian framework. They were able to constrain the radius of a neutron star with mass $M = 1.4\, M_\odot$ to values around $11 - 12$ km.

1.4 Spin periods

By identifying the periodicity of the arrival time of pulses of radiation with the neutron star spin periods implies that these celestial bodies rotate at very high speeds. From the observed population their spin periods range from a few milliseconds to a few tens of seconds (see Figure 1.12). The majority of these measurements come from radio timing analysis which has also unveiled that their spin periods increase with time at rates between around $10^{-20}\,\mathrm{s\,s^{-1}}$ and $10^{-10}\,\mathrm{s\,s^{-1}}$. This suggests that generally neutron stars spin down over time and they should be born with even shorter rotation periods.

Several works involving analyses of young supernova remnants (e.g. Chevalier, 2005), population synthesis studies (e.g. Popov and Turolla, 2012; Faucher-Giguère and Kaspi, 2006; Gullón et al., 2014) and supernova simulations (e.g. Ott et al., 2006; Janka, Wongwathanarat, and Kramer, 2022), have suggested that the birth spin periods of neutron stars should range between a few tens of milliseconds up to a few hundreds of milliseconds. However, the origin of such fast rotation is still unclear.

A first insight into how neutron stars could achieve these fast rotation can be gained by considering the angular momentum conservation during the collapse of the iron core. For an order of magnitude estimation we consider a core with mass $M_{\mathrm{core}} \sim 1.4\, M_\odot$, radius comparable to our Sun $R_{\mathrm{core}} \sim 10^{11}$ cm and spin period $P \sim 1$ day. From the conservation of angular momentum, we get that:

$$M_{\mathrm{core}}R_{\mathrm{core}}^2\omega_{\mathrm{core}} \sim M_{\mathrm{NS}}R_{\mathrm{NS}}^2\omega_{\mathrm{NS}}, \tag{1.3}$$

where $\omega = 2\pi/P$ is the spin angular velocity and we assumed mass conservation during the core collapse, i.e., $M_{\mathrm{core}} \sim M_{\mathrm{NS}}$. During the collapse, as the radius shrinks to $R_{\mathrm{NS}} \sim 10^6$ cm reducing the moment of inertia by a factor $\sim 10^{10}$, the angular velocity should increase by the same amount. As a consequence, newborn neutron stars should spin with periods of the order of milliseconds.

However angular momentum conservation is not the only ingredient to explain the neutron star's fast rotation. Various mechanisms acting during and after the supernova explosion could contribute to spinning up or spinning down the collapsing

core and the newly formed neutron star. During the core collapse and the development of the supernova, hydrodynamical instabilities and turbulent motion of mass could cause asymmetric mass ejection in the supernova. This asymmetry imparts accelerating forces on the core over timescales from several seconds to hours. In this way, the newly formed neutron star is displaced from the centre of the collapse and receives a kick that can be of several hundreds of kilometers per second (see, e.g., Hobbs et al., 2005, but also Chapter 5). As the neutron star moves in the supernova ejecta it can accrete a fraction of the matter that has remained gravitationally bound. This anisotropic accretion of mass could transfer a relevant amount of angular momentum to the compact remnant after the development of the supernova Janka, Wongwathanarat, and Kramer (2022) and Coleman and Burrows (2022). This scenario suggests also a possible link between the origin of the fast neutron stars' proper motion and their spin period.

1.5   Magnetic fields

To our knowledge neutron stars are also the strongest magnets in the Universe. A direct measurement of the magnetic field strength is not straightforward though. One possibility is through the detection of cyclotron lines, that generally appear as absorption features in the UV and X-ray spectra of systems containing highly magnetised neutron stars. Almost all the observations of cyclotron absorption features come from neutron stars accreting matter from a stellar companion, also called X-ray binary pulsars. These features are produced as the photons produced by the hot column of plasma accreted onto the neutron star's magnetic poles are resonantly scattered by charged particles gyrating around the magnetic field lines (see for example Staubert et al., 2019, for a review).

However, generally, neutron star magnetic fields are indirectly estimated from their rotational properties. The commonly adopted explanation for this rotational energy loss is that the neutron star experiences an electromagnetic torque generated by a very powerful dipolar magnetic field (see Section 1.8). With this assumption, from the measurements of the spin period and spin period derivative one can infer dipolar magnetic field values that range between 10^8 G to almost 10^{15} G.

How such strong magnetic fields are generated is still a matter of debate. One idea is that the neutron star inherits the magnetic field of the parent OB star. In this fossil field hypothesis, during the core collapse the magnetic field gets amplified due to magnetic flux freezing (Ferrario and Wickramasinghe, 2006; Hu and Lou, 2009). Ordinary OB stars have a magnetic field of the order of $B \sim 10^2$ G with some of them exceeding 10^3 G. From the conservation of the magnetic flux, we have that:

$$B_{\mathrm{core}} R_{\mathrm{core}}^2 \sim B_{\mathrm{NS}} R_{\mathrm{NS}}^2. \qquad (1.4)$$

As the core radius contracts by a factor $\sim 10^5$, the magnetic field should be amplified by a factor $\sim 10^{10}$. This translates to a magnetic field of the order of $B_{\mathrm{NS}} \sim 10^{12}$ G for a newborn neutron star, in line with current estimates. A second hypothesis is that a strong magnetic field originates dynamically from a dynamo process. During the formation of the proto-neutron star, the magnetic field could be amplified by electric currents generated from rapid rotation or turbulent convective motions of conductive material inside the proto-neutron star a few seconds after its formation (see e.g. Burrows, 1987; Burrows and Lattimer, 1988; Duncan and Thompson, 1992; White et al., 2022). Probably both effects play a role in the creation of such intense magnetic fields.
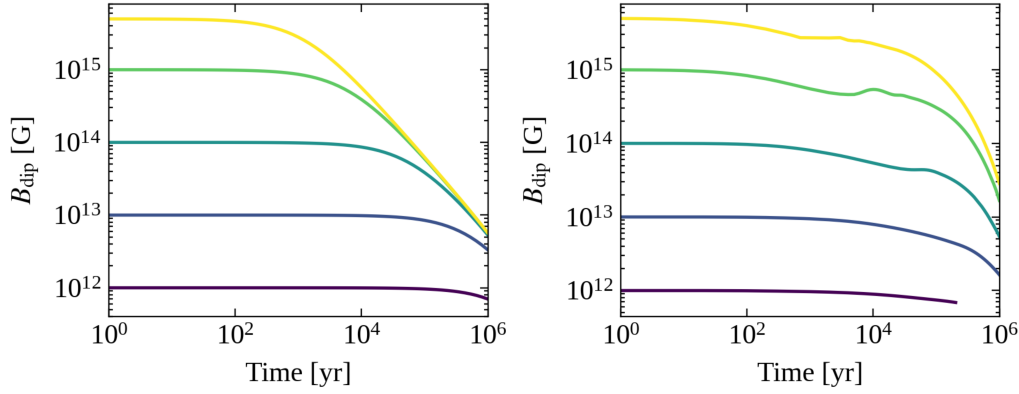
FIGURE 1.2: Examples of dipolar magnetic field evolution curves. The analytical solution in Equation (1.15) (*left panel*) is compared to the numerical solution from 2D magneto-thermal simulations (Viganò et al., 2021) (*right panel*). In the right panel, the purple curve corresponding to an initial magnetic field of 10^{12} G appears truncated due to code constraints in modelling the micro-physical behaviour of the neutron-star crust at low temperatures (see also Appendix C for more details).

1.6 Magnetic field evolution

Much effort has been put into modelling the magnetic field geometry and evolution inside a neutron star. Soon after birth, a solid crust of about 1 km thickness is formed. At the densities of the crust ($10^{10} - 10^{14}$ g cm^{-3}), matter is pressure-ionised. Atoms are so densely packed together that electrons are not bound to a single nucleus but feel the electric fields of all the neighbouring nuclei. As a consequence, while the ions in the lattice or the pasta phase layer (see Section 1.3) have very restricted mobility, electrons are free to flow constituting the main source of conductivity. Therefore describing the magnetic field evolution in the crust is more straightforward as the electrons are the only charged particles affecting it and the microphysical properties of matter are relatively well known. On the other hand, describing the field geometry and evolution in the liquid core is more complex. Due to the extreme conditions prevalent within the core, the exact composition and the properties of matter are still uncertain. In the fluid core, besides electrons, also neutrons, protons, muons and possibly more exotic particles could influence the magnetic field evolution. In particular it is believed that neutron superfluidity and proton superconductivity can play a critical role for the internal dynamics (see for example Chamel, 2017; Wood and Graber, 2022). In general, for the core, it is very challenging to accurately incorporate realistic microphysics and estimate the impact of these effects on the magnetic field's evolution (Viganò et al., 2021). Here we will focus instead on the evolution of the magnetic field which is confined in the neutron star crust. In general, the equation governing the magnetic field evolution is Maxwell's induction equation:

$$\frac{\partial \boldsymbol{B}}{\partial t} = -c \left(\boldsymbol{\nabla} \times \boldsymbol{E} \right). \tag{1.5}$$

For a conducting fluid of electrons, moving with velocity $\boldsymbol{v}_e$ and in the presence of a magnetic field $\boldsymbol{B}$, the electric field $\boldsymbol{E}$ is related to the current density $\boldsymbol{j}$ through the

generalised Ohm's law:

$$E = \frac{j}{\sigma_e} - \frac{v_e}{c} \times B, \tag{1.6}$$

where σ_e is the electrical conductivity that can be defined in terms of the mean free time between collisions or relaxation time τ_e (Longair, 2011):

$$\sigma_e = \frac{e^2 n_e \tau_e}{m_e}. \tag{1.7}$$

The electron velocity is also related to the current and electron density n_e trough:

$$v_e = -\frac{j}{e n_e}. \tag{1.8}$$

We can also assume that the timescale for collisions inside the plasma is much shorter than the timescale of the variation of the electromagnetic field so that we can neglect the displacement current term in Ampere's law and express the current density as:

$$j = -\frac{c}{4\pi} \left(\nabla \times B \right). \tag{1.9}$$

Therefore by combining the four equations above, the induction equation assumes the form (Pons and Geppert, 2007; Viganò, Pons, and Miralles, 2012):

$$\frac{\partial B}{\partial t} = -\nabla \times \left(\frac{c^2}{4\pi\sigma_e} \nabla \times B + \frac{c}{4\pi e n_e} (\nabla \times B) \times B \right). \tag{1.10}$$

The first term in the parentheses describes the Ohmic dissipation. It is responsible for converting magnetic energy into heat according to $Q = j^2/\sigma_e$ which is then deposited in the crust. In a neutron star crust, the conductivity σ_e is dominated by the electronic transport and depends on the crustal temperature, the electron density, and the impurity concentration. To quantify the impurity content of the lattice we define the impurity parameter as the mean quadratic deviation of the atomic number Z:

$$\mathcal{Q}_{\mathrm{imp}} = \sum_i Y_i \left(Z_i^2 - \langle Z^2 \rangle \right), \tag{1.11}$$

where Y_i is the relative abundance of the nuclide with $Z_i e$ charge, and $\langle Z^2 \rangle$ is the average squared atomic number in the lattice (see Potekhin, Pons, and Page, 2015, and reference therein). In general, for higher temperatures, and for larger impurity concentration the conductivity is lower as there is more scattering between the electrons and the lattice. As a result the magnetic field evolution is highly coupled to the temperature evolution. The higher the magnetic field, the greater the dissipation and the heat deposited into the crust. In this way the stellar surface is kept hot for longer times. This in turn affects the microphysical properties of the crust and the electrical conductivity so that the heat deposition itself influences how the magnetic field dissipates. Furthermore, the electron density varies over about four orders of magnitude in the crust. As a consequence, the electric conductivity varies by many orders of magnitude across the neutron star crust as well. What drives the evolution is usually the resistivity of the region where currents are predominantly placed. For example, the dissipation will be greater if currents are mainly circulating in the pasta layer of the crust. Here the transport properties are uncertain but the irregularities in

the nuclear matter shapes and charge distribution should enhance the resistivity and therefore the magnetic field dissipation (Horowitz et al., 2015; Nandi and Schramm, 2018).

The second term in Equation 1.10 is a non-linear term which strongly depends on the magnetic field strength and is related to the Hall drift. It does not cause dissipation, but it is responsible for the drift and rearrangement of the electric currents inside the crust and the redistribution of the magnetic energy on smaller scales. By affecting the distribution and scale of the electric currents the Hall effect is responsible of enhancing the Ohmic dissipation. As this term depends only on the degenerate electron density it should be insensitive to temperature (see also Cumming, Arras, and Zweibel, 2004).

If one is concerned only with the evolution of the magnitude of the magnetic field in time, the induction equation can be rewritten in an approximated form as (see Aguilera, Pons, and Miralles, 2008a; Aguilera, Pons, and Miralles, 2008b):

$$
\begin{aligned}
\frac{\mathrm{d}B}{\mathrm{d}t} &\simeq -\frac{c^2 B}{4\pi\sigma L^2} - \frac{cB^2}{4\pi e n_e L^2} \\
&= -\frac{B}{\tau_{\mathrm{Ohm}}} - \frac{B^2}{B_0 \tau_{\mathrm{Hall},0}},
\end{aligned}
\tag{1.12}
$$

where the spatial derivatives have been replaced by a factor $1/L$ where L is a typical length scale over which the physical quantities B, n_e and σ_e change. We can also define two typical timescales, one for the Ohmic dissipation, and the other for the Hall drift:

$$
\begin{aligned}
\tau_{\mathrm{Hall},0} &= \frac{4\pi e n_e L^2}{c B_0} \\
&\simeq 6.4 \times 10^4 \,\mathrm{yr} \left(\frac{n_e}{10^{35}\,\mathrm{cm}^{-3}}\right) \left(\frac{L}{1\,\mathrm{km}}\right)^2 \left(\frac{B_0}{10^{14}\,\mathrm{G}}\right)^{-1},
\end{aligned}
\tag{1.13}
$$

$$
\begin{aligned}
\tau_{\mathrm{Ohm}} &= \frac{4\pi\sigma L^2}{c^2} \\
&\simeq 4.4 \times 10^6 \,\mathrm{yr} \left(\frac{\sigma}{10^{24}\,\mathrm{s}^{-1}}\right) \left(\frac{L}{1\,\mathrm{km}}\right)^2,
\end{aligned}
\tag{1.14}
$$

where the numerical values are taken from Cumming, Arras, and Zweibel (2004). Note that the Hall timescale defined here is inversely proportional to the initial magnetic field magnitude B_0. By neglecting the influence of the temperature we find the following analytical solution for Equation (1.12):

$$
B(t) = B_0 \frac{e^{-t/\tau_{\mathrm{Ohm}}}}{1 + \frac{\tau_{\mathrm{Ohm}}}{\tau_{\mathrm{Hall},0}} \left[1 - e^{-t/\tau_{\mathrm{Ohm}}}\right]}.
\tag{1.15}
$$

Even if it is a very crude approximation, this equation captures a first stage that is characterised by rapid (non-exponential) decay regulated by the Hall timescale $\tau_{\mathrm{Hall},0}$, and a second stage that is characterised by exponential decay due to Ohmic dissipation and regulated by the timescale τ_{Ohm} (see left panel in Figure 1.2).

To model the magnetic field evolution and decay in neutron stars consistently, a complete magneto-thermal simulation is required that takes into account the coupled magnetic field and temperature evolution. Moreover inside the neutron star general relativistic effects become relevant. Therefore the induction equation (1.10) has to be corrected to take into account the curvature of space-time. Generally the

shwarzschild static metric is sufficient as any deformation induced by rotation and the magnetic fields are negligible. In recent years numerical codes able to simulate the magnetic field and temperature evolution of neutron stars have been developed. In particular large efforts have been put into 2D axis-symmetric simulations (see for example Aguilera, Pons, and Miralles, 2008b; Pons, Miralles, and Geppert, 2009; Viganò, Pons, and Miralles, 2012; Viganò et al., 2021). More recently the more complex but more realistic 3D case is being explored (see Gourgouliatos, Wood, and Hollerbach, 2016; De Grandis et al., 2020; Dehman et al., 2023). The right panel of Figure 1.2 shows the magnetic field evolution curves predicted by 2D magneto-thermal simulations (Viganò et al., 2021).

1.7 The magnetosphere

The magnetosphere of a neutron star is a complex and dynamic region filled with magnetised plasma where strong magnetic fields interact with the surrounding environment. It envelopes the entire star extending from the stellar surface up to several thousands of kilometers. Unlike Earth's magnetosphere, which is primarily influenced by the interaction with the solar wind, neutron star magnetospheres can encounter various external influences, including the interstellar medium, accretion disks, and even companion stars in binary systems. Deutsch (1955) was the first to study analytically the structure of the electromagnetic field surrounding an isolated rotating magnetised star with a misaligned magnetic dipole moment in a vacuum. Later, Goldreich and Julian (1969) described the pulsar electrodynamics in the simplest case of a rotating magnetic dipole, aligned with the rotational axis, surrounded by a charge-separated plasma. In general analytical solutions for the magnetospheric structure are possible only in very simplified and ideal cases. To fully study the overall geometry of the electromagnetic fields and how the surrounding plasma interacts with and influences them, numerical simulations are required. Contopoulos, Kazanas, and Fendt (1999) were the first to present a numerical, realistic configuration for an aligned rotator in the force-free approximation (see below). His work was the first step towards more realistic 3D simulations that were able to quantitatively consider the effects of the misalignment between rotation and magnetic axes (see Figure 1.4 and Spitkovsky, 2006; Philippov, Tchekhovskoy, and Li, 2014). Understanding the structure and dynamics of neutron star magnetospheres is crucial for comprehending a wide range of phenomena associated with these celestial objects. In particular, the intricate interplay between the magnetic field and the highly energetic particles in the surrounding plasma leads to the generation of electromagnetic radiation across different wavelengths from radio to gamma rays. For more detailed reviews about pulsar magnetosphere see for example Pétri (2016) and Pétri (2020).

   To gain an insight into the magnetosphere's structure we consider a neutron star initially surrounded by vacuum spinning with angular frequency $\omega = 2\pi/P$. The neutron star is endowed with a strong magnetic field. However to define this magnetic field we need to make a distinction between two reference frames: an inertial reference frame of an external observer and a reference frame co-rotating with the star. For simplicity we consider the latter to be sufficiently small compared to the star's size and not so far from the axis of rotation as to be considered locally inertial. The position of the co-rotating frame can be described by a distance vector, r, originating r at the centre of the neutron star. With these considerations, the electromagnetic fields as measured by the two observers can be linked by the Lorentz
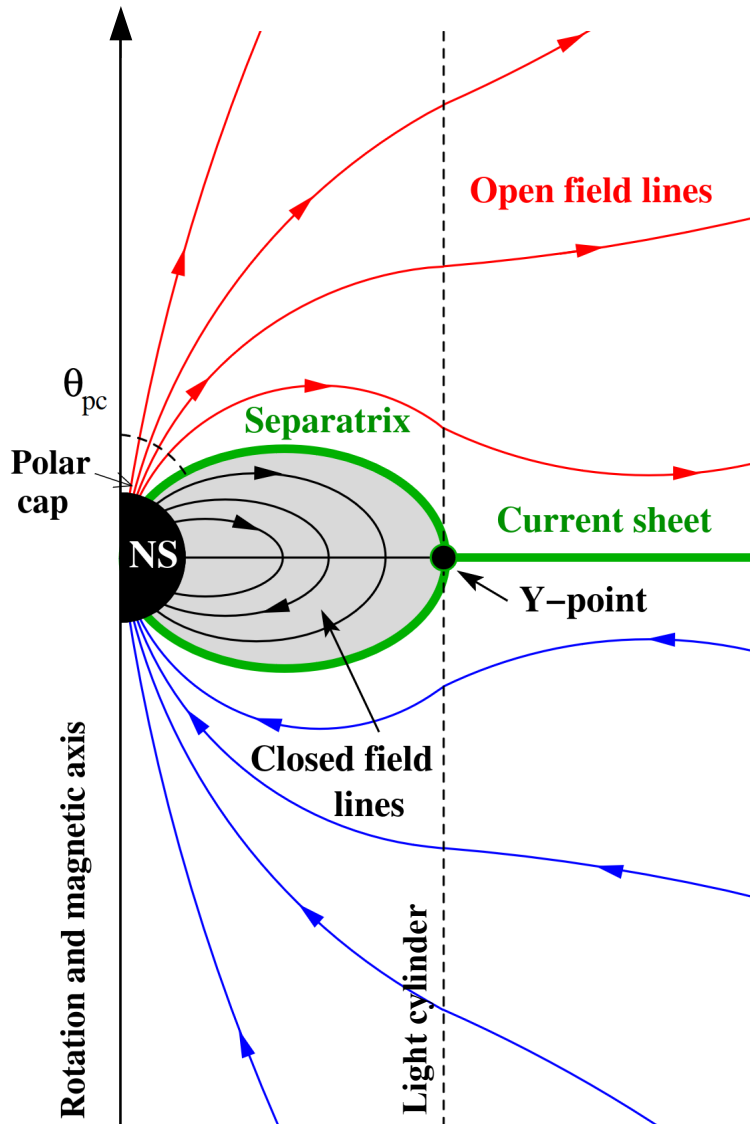
FIGURE 1.3: Sketch of the ideal force-free magnetosphere of the aligned pulsar. The main elements are: (i) The closed field line region (grey, and black field lines) lying between the star surface and the light cylinder. This zone is dead and does not participate to the pulsar activity. (ii) The open field line region (red and blue field lines) extending beyond the light cylinder. As particles are accelerated along the open field line region they emit electromagnetic radiation. (iii) The equatorial current sheet (green) between the opposite magnetic fluxes in the open field line region. It splits at the light cylinder into two separatrix current sheets that go around the closed zone, between the last open and the first closed field lines (see Cerutti and Beloborodov, 2017, for further details) (Figure adapted from Cerutti and Beloborodov, 2017).
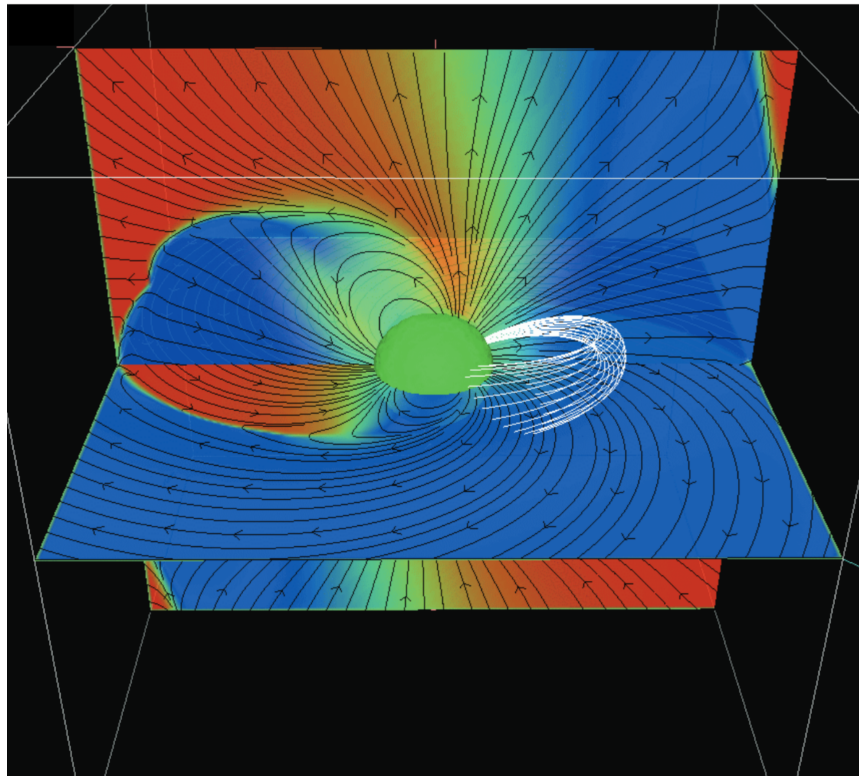
FIGURE 1.4: Slices through the simulated force-free magnetosphere of a neutron star with inclination angle of 60° between the magnetic dipole moment and the rotation axis. The field lines in the horizontal and vertical slices are shown; colour represents the magnetic field component perpendicular to the slices. A sample three-dimensional magnetic flux tube is traced in white (Figure taken from Spitkovsky, 2006).

transformations (Jackson, 1998):

$$B = \gamma_{\text{rot}} \left(B' + \beta_{\text{rot}} \times E' \right) - \frac{\gamma_{\text{rot}}^2}{\gamma_{\text{rot}} + 1} \beta_{\text{rot}} (\beta_{\text{rot}} \cdot B'), \tag{1.16}$$

$$E = \gamma_{\text{rot}} \left(E' - \beta_{\text{rot}} \times B' \right) - \frac{\gamma_{\text{rot}}^2}{\gamma_{\text{rot}} + 1} \beta_{\text{rot}} (\beta_{\text{rot}} \cdot E'), \tag{1.17}$$

where E and B are the fields measured by an external observer, while E' and B' are the fields measured in the co-rotating frame, $\beta_{\text{rot}} = (\omega \times r)/c$ is the ratio of the rotational velocity to the speed of light c and $\gamma_{\text{rot}} = 1/(1 - \beta_{\text{rot}}^2)^{1/2}$ is the corresponding Lorentz factor. If we consider the neutron star interior or we stay in the proximity of its surface, for typical observed spin periods we have that $\beta_{\text{rot}} \equiv |\beta_{\text{rot}}| \ll 1$. Therefore the transformations above assume the simplified form:

$$B \sim \left(B' + \beta_{\text{rot}} \times E' \right), \tag{1.18}$$

$$E \sim \left(E' - \beta_{\text{rot}} \times B' \right), \tag{1.19}$$

where we neglected second order terms in β_{rot} and considered $\gamma_{\text{rot}} \sim 1$.

Little is known about the magnetic field geometry inside the neutron star. In the exterior, the magnetic field is typically described by a combination of multipolar components. As the higher-order multipoles decay faster with distance from the stellar surface, the dipolar component probably dominates the magnetic field geometry at a large scale. For the sake of simplicity we consider that in the neutron star's co-rotating frame the magnetic field is purely dipolar. In this case the formula that relates the magnetic moment μ with the strength of the magnetic field B' is given by:

$$B'(r) = \frac{3(r \cdot \mu)r}{r^5} - \frac{\mu}{r^3}. \tag{1.20}$$

We further simplify by assuming a configuration with the magnetic moment μ aligned with the rotation axis of the star. We will call this an aligned rotator. In spherical coordinates (r, θ, ϕ) with θ being the angle computed from the magnetic moment axis (or spin axis in the aligned case) and ϕ the azimuthal angle, the dipolar magnetic field can be expressed as:

$$B'(r, \theta) = \frac{\mu}{r^3} \begin{bmatrix} 2\cos\theta \\ \sin\theta \\ 0 \end{bmatrix}. \tag{1.21}$$

Due to the abundance of free charges in its interior, as a first approximation a neutron star can be considered as a perfect conductor. Therefore in the co-rotating frame, the internal electric field has to vanish so that in the neutron star interior the following equations are valid:

$$B \sim B', \tag{1.22}$$

$$E \sim -\beta_{\text{rot}} \times B' = -\frac{\omega \times r}{c} \times B'. \tag{1.23}$$

In other words, a perfect conductor rotating in a magnetic field behaves as a unipolar inductor converting the rotational energy into electrical energy. As seen by an inertial external observer the rotation of the star induces an electric field inside the star and the free charges inside rearrange themselves under its action. In particular, on the surface of the star particles of opposite charges are driven to the polar and equatorial regions respectively according to their signs. This creates a quadrupolar electric potential, V, on the surface that can be derived from $\boldsymbol{E} = -\boldsymbol{\nabla}V$.

In particular, for an aligned rotator, from Equation (1.23), the θ component of the electric field computed at the surface is given by:

$$E_\theta = -\frac{\omega R_{\mathrm{NS}} \sin\theta B_r}{c}. \tag{1.24}$$

This determines the electrostatic potential drop with latitude on the stellar surface that can be computed considering that $E_\theta = 1/r(\partial V/\partial\theta)$ and integrating over the coordinate θ. By using the radial component of the surface dipole magnetic field in Equation (1.21), $B_r = 2\mu\cos\theta/R_{\mathrm{NS}}^3$, we find (Backus, 1956; Goldreich and Julian, 1969; Cerutti and Beloborodov, 2017):

$$V(R_{\mathrm{NS}},\theta) - V(R_{\mathrm{NS}},0) = \int_0^\theta E_\theta R_{\mathrm{NS}}\mathrm{d}\theta = \frac{\omega\mu}{2R_{\mathrm{NS}}c}\sin^2\theta, \tag{1.25}$$

where $V(R_{\mathrm{NS}},0) \equiv V_{\mathrm{p}}$ is the value of the electric potential at the pole. This charge redistribution induces an external quadrupolar electric field with a strong component parallel to the magnetic field lines. As shown by Goldreich and Julian (1969) this component of the electric field is strong enough to overcome gravity and pull out charged particles from the neutron star surface unless strong cohesive forces prevent them from escaping. This happens primarily for the electrons and protons as, in the presence of a strong magnetic field, heavier ions in the surface layer can rearrange themselves to form a dense and strongly bound molecular lattice (see Ruderman and Sutherland, 1975, and reference therein). Even if considered as a simplified and ideal case, the result from Goldreich and Julian (1969) implies that a vacuum region around a neutron star cannot be maintained and a magnetosphere filled with conducting and co-rotating plasma has to surround the star.

As charges are extracted from the surface they are accelerated up to relativistic speeds along the magnetic field lines by the induced electric field. By moving along curved field lines these charges produce gamma-ray photons by curvature radiation (Ruderman and Sutherland, 1975) or the inverse-Compton process by up-scattering lower energy photons (Daugherty and Harding, 1986). In the presence of strong magnetic fields gamma-ray photons can transform into electron-positron pairs, if the photon energy exceeds the threshold $E_\gamma \geq 2m_e c^2$ where m_e is the electron mass. This generation of new particles then produces new photons and pairs leading to a pair cascade that populates the magnetosphere with co-rotating highly conducting plasma. Assuming a continuous supply of charges, as the magnetosphere is replenished with perfectly conducting plasma (predominantly formed by $e^- e^+$ pairs) the electric field component along the magnetic field lines gets screened so that in a steady state one expects the condition $\boldsymbol{E}\cdot\boldsymbol{B} = 0$ to be satisfied. This prevents any further acceleration of particles along the field lines so that these can be considered equipotential lines. In this ideal case, called force-free limit, the magnetosphere is sufficiently populated with plasma for the conductivity in the medium to become infinite. Moreover, the electromagnetic field dominates the dynamics of the magnetosphere by several orders of magnitude with respect to pressure, gravity and

inertia. The unipolar induction mechanism on the other hand induces an electric field perpendicular to the magnetic field lines that lead to charge separation also in the magnetosphere.

However, the flow of plasma in the magnetosphere has a retroactive effect on the electromagnetic field structure around the neutron star. As we move further away from the neutron star surface, the approximations and the electromagnetic field configuration outlined above break down. As the rotational velocity approaches the speed of light plasma co-rotation becomes impossible and a co-rotating frame cannot be defined anymore. We define an imaginary cylindrical surface called light cylinder, with the axis parallel to the rotation axis of the star. The radius of the light cylinder corresponds to the distance where the co-rotating speed is equal to the speed of light:

$$r_{lc} = \frac{c}{\omega} \simeq 4.77 \times 10^9 \, \text{cm} \left(\frac{P}{1\,\text{s}} \right). \tag{1.26}$$

Near the light cylinder plasma inertia due to the approaching of the speed of light becomes relevant and the plasma is not able to stay in co-rotation. In particular, the plasma drags the magnetic field causing the field lines to sweep back in the azimuthal direction and spiral around the star (see Figure 1.4). This produces a toroidal magnetic field component. Furthermore, the co-rotating charges and these inertial effects create currents which modify the magnetic field configuration. These magnetic perturbations tend to repel field lines and as a result, the field lines near the light cylinder inflate to infinity and open up. In the vicinity of the light cylinder, the force-free condition is violated and dissipation effects and particle acceleration create electromagnetic emission detectable on Earth. The magnetosphere can then be divided into two regions with different properties (see Figure 1.3). The magnetic field lines that do not cross the light cylinder form the so-called closed magnetosphere (grey region in Figure 1.3). Here the plasma is trapped and the condition $E \cdot B = 0$ is valid everywhere. This region should therefore be electromagnetically inactive as particle acceleration is suppressed. The field lines that cross the light cylinder and are inflated to infinity form the open field line region (red and blue lines in Figure 1.3). Here charges are free to stream along the magnetic field lines and escape to infinity. Therefore here the condition $E \cdot B = 0$ can be violated and gaps with low-density plasma can be created where charged particles are accelerated. The open field line region is therefore active and it is the region where electromagnetic emission from pulsars is believed to be produced.

1.8 Dipolar spin-down evolution

The observation of the secular increase of spin periods for the detected neutron stars implies that they are losing rotational energy over time. Assuming that neutron stars are solid bodies the loss of rotational energy can be computed as:

$$\frac{dE_{rot}}{dt} = \frac{d}{dt} \left(\frac{1}{2} I_{NS} \omega^2 \right) = I_{NS} \omega \frac{d\omega}{dt} = -4\pi^2 I_{NS} \frac{dP}{dt} \frac{1}{P^3}$$

$$\simeq -5 \times 10^{31} \, \text{erg s}^{-1} \left(\frac{\dot{P}}{10^{-15}\,\text{s s}^{-1}} \right) \left(\frac{P}{1\,\text{s}} \right)^{-3}, \tag{1.27}$$

where we define the spin-down rate $\dot{P} \equiv dP/dt$ and assume a neutron star to be a perfect sphere of mass $M_{NS} = 1.4\, M_{\odot}$ and radius $R_{NS} = 11\,\text{km}$ and with moment of inertia $I_{NS} \simeq 2/5 M_{NS} R_{NS}^2 \simeq 1.4 \times 10^{45}\,\text{g}\,\text{cm}^2$ which is constant in time.

If a neutron star can be simply pictured as a rotating magnetic dipole with moment vector $\boldsymbol{\mu}$ misaligned by an angle χ with respect to the rotation axis, the spin-down can be explained using classical electrodynamics arguments (see for example Shapiro and Teukolsky, 1983). In classical electrodynamics a rotating misaligned magnetic dipole emits electromagnetic energy (see Jackson, 1998). In analogy with the Larmor formula for an accelerated charge, a time-varying magnetic dipole radiates away a total power integrated over the solid angle equal to:

$$\frac{dE}{dt} = \frac{2}{3c^3}\left|\frac{d^2\boldsymbol{\mu}}{dt^2}\right|^2. \tag{1.28}$$

For a misaligned magnetic dipole with an inclination angle χ, rotating at an angular frequency ω, we can write the magnetic dipole moment components as a function of time t in Cartesian coordinates. We assume rotation about the z axis:

$$\boldsymbol{\mu} = \mu \begin{bmatrix} \sin\chi\cos\omega t \\ \sin\chi\sin\omega t \\ \cos\chi \end{bmatrix}. \tag{1.29}$$

where $\mu = |\boldsymbol{\mu}|$. If we compute the second time derivative for this vector we find:

$$\frac{d^2\boldsymbol{\mu}}{dt^2} = \mu \begin{bmatrix} -\omega^2\sin\chi\cos\omega t \\ -\omega^2\sin\chi\sin\omega t \\ 0 \end{bmatrix}. \tag{1.30}$$

Here we have assumed that the spin frequency, the magnetic moment and the inclination angle change on timescales much longer than a rotation period so that they can be assumed as constants in computing these derivatives, i.e., they do not explicitly depend on the time t. By computing the squared module of the vector above and substituting it into Equation (1.28) we get an emitted power:

$$\frac{dE}{dt} = \frac{2\mu^2}{3c^3}\omega^4\sin^2\chi. \tag{1.31}$$

By assuming that this electromagnetic energy is emitted at the expense of the rotational energy we can equate Equation (1.27) with Equation (1.31) to obtain an equation that describes how the spin frequency changes in time:

$$\frac{d\omega}{dt} = -\frac{2\mu^2}{3c^3 I_{NS}}\omega^3\sin^2\chi. \tag{1.32}$$

We find an analogous result by considering a more sophisticated calculation. By picturing the neutron star as a spherical conductor rotating in an electromagnetic field in vacuum it is possible to compute the net torque acting on it. In particular the electric and magnetic fields generate forces on the charges in the conductor which in turn generate a net torque acting on the conductor surface. This torque, $\boldsymbol{M}$, can be estimated by the following integral computed over the neutron star surface S (see

Michel and Goldwire, 1970; Philippov, Tchekhovskoy, and Li, 2014):

$$M = \int_S n \cdot (r \times T) \, dS,$$ (1.33)

where r is the radial vector originating in the neutron star centre, T is the Maxwell stress tensor and n is the unit vector perpendicular to the surface. By considering the expressions for the electromagnetic fields near the stellar surface computed by Deutsch (1955) for a rotating star characterised as a misaligned magnetic dipole in vacuum, Michel and Goldwire (1970) computed the non-vanishing torque components. The first one is antiparallel to the spin frequency vector and causes the neutron star to spin down:

$$M_\parallel = -\frac{2\mu^2}{3c^3}\omega^3 \sin^2 \chi.$$ (1.34)

The second one is perpendicular to the spin frequency vector and acts in the plane defined by the two vectors ω and μ. While it does not change the spin frequency, it is responsible for the alignment of the magnetic axis with the rotational axis:

$$M_\perp = -\frac{2\mu^2}{3c^3}\omega^3 \sin\chi \cos\chi.$$ (1.35)

A third component is perpendicular to the $\omega - \mu$ plane and causes precession of the spin frequency vector around the magnetic axis (Philippov, Tchekhovskoy, and Li, 2014). In the following, we will not take into account this precession since it does not affect the evolution of the spin period and inclination angle. The evolution of the angular momentum vector, L, of the neutron star is then given by:

$$\frac{dL}{dt} = I_{NS}\frac{d\omega}{dt} = M,$$ (1.36)

which gives the following coupled differential equations:

$$\frac{d\omega}{dt} = -\frac{2\mu^2}{3c^3 I_{NS}}\omega^3 \sin^2\chi,$$ (1.37)

$$\frac{d\chi}{dt} = -\frac{2\mu^2}{3c^3 I_{NS}}\omega^2 \sin\chi \cos\chi.$$ (1.38)

Note that the first equation is the same as Equation (1.32). In general, neutron stars evolves towards longer periods and a configuration where the magnetic dipole axis is aligned with the rotation axis. Therefore rotational evolution tends towards a state where the energy losses due to the torques are minimised as can be seen from Equation (1.31). However, these equations are valid in the case of a vacuum magnetosphere which as we have seen in Section 1.7 is not the most realistic scenario.

For plasma-filled magnetospheres, an analytical solution of the torques can not be derived and simulations are required due to the non-linearity of the problem. Through 3D simulations of force-free plasma-filled magnetospheres Spitkovsky (2006) found that the emitted electromagnetic luminosity can be well described by the following relation:

$$\left(\frac{dE}{dt}\right)_{force-free} = \frac{\mu^2}{c^3}\omega^4(\kappa_0 + \kappa_1 \sin^2\chi).$$ (1.39)

The rotational evolution of a neutron star can then be approximated by the two following equations:

$$\frac{\mathrm{d}\omega}{\mathrm{d}t} = -\frac{\mu^2}{c^3 I_{\mathrm{NS}}}\omega^3\left(\kappa_0 + \kappa_1 \sin^2\chi\right),\tag{1.40}$$

$$\frac{\mathrm{d}\chi}{\mathrm{d}t} = -\frac{\mu^2}{c^3 I_{\mathrm{NS}}}\omega^2\left(\kappa_2 \sin\chi\cos\chi\right),\tag{1.41}$$

and substituting $\omega = 2\pi/P$ we get:

$$\frac{\mathrm{d}P}{\mathrm{d}t} = \frac{4\pi^2\mu^2}{c^3 I_{\mathrm{NS}}P}\left(\kappa_0 + \kappa_1 \sin^2\chi\right),\tag{1.42}$$

$$\frac{\mathrm{d}\chi}{\mathrm{d}t} = -\frac{4\pi^2\mu^2}{c^3 I_{\mathrm{NS}}P^2}\left(\kappa_2 \sin\chi\cos\chi\right),\tag{1.43}$$

The parameters κ_0, κ_1 and κ_2 are derived by fitting simulation results and for force-free magnetosphere they take the values $\kappa_0 \simeq \kappa_1 \simeq \kappa_2 \simeq 1$ (Philippov, Tchekhovskoy, and Li, 2014). Compared to the vacuum solution (where $\kappa_0 = 0$, $\kappa_1 = 2/3$ and $\kappa_2 = 1$) the electromagnetic torque and the energy loss does not vanish when the magnetic dipole axis aligns with the rotation axis and the torque exerted on the star can be a factor ~ 2 stronger.

1.8.1   Dipole magnetic field estimate

By considering the magnetic dipole breaking model outlined above we can estimate the dipolar magnetic field of a neutron star from its timing properties P and $\dot{P}$. If we consider the magnetic field strength at the magnetic pole of the neutron star at $r = R_{\mathrm{NS}}$, from Equation (1.20) we get the expression:

$$B_{\mathrm{p}} = \frac{2\mu}{R_{\mathrm{NS}}^3} \implies \mu = \frac{B_{\mathrm{p}}R_{\mathrm{NS}}^3}{2}.\tag{1.44}$$

If we consider the magnetic field strength at the magnetic equator of the neutron star we obtain that it is half of the value at the magnetic pole:

$$B_{\mathrm{e}} = \frac{\mu}{R_{\mathrm{NS}}^3} \implies \mu = B_{\mathrm{e}}R_{\mathrm{NS}}^3.\tag{1.45}$$

We can substitute Equation (1.44) or Equation (1.45) into Equation (1.42) to obtain a relation between the observed quantities P and $\dot{P}$ and the magnetic field strength at the poles or at the equator. In the literature, the vacuum magnetosphere model is often assumed even if less realistic. This usually implies an overestimation of the magnetic field strength by a factor of ~ 2 with respect to the force-free magnetosphere solution. By assuming for the sake of simplicity an inclination angle $\chi = \pi/2$ (since it is usually unknown) we obtain a formula to roughly estimate the strength of the magnetic field at the neutron star poles or equator as a function of the timing properties:

$$B_{\mathrm{p}} = \left(\frac{3c^3}{2\pi^2}\frac{I_{\mathrm{NS}}P\dot{P}}{R_{\mathrm{NS}}^6}\right)^{1/2} \simeq 2\times10^{12}\,\mathrm{G}\left(\frac{P}{1\,\mathrm{s}}\right)^{1/2}\left(\frac{\dot{P}}{10^{-15}\,\mathrm{s}\,\mathrm{s}^{-1}}\right)^{1/2},\tag{1.46}$$

where we consider a neutron star moment of inertia $I_{\mathrm{NS}} \simeq 10^{45}\,\mathrm{g\,cm^2}$ and radius $R_{\mathrm{NS}} \simeq 10^6$ cm.

1.8.2 Characteristic age estimate

By assuming that both the magnetic field and the inclination angle stay constant during the neutron star's life, one can integrate of Equation (1.40) in time to obtain a characteristic age:

$$\tau_{\mathrm{c}} = -\frac{\omega^3}{2\dot{\omega}} \left(\frac{1}{\omega^2} - \frac{1}{\omega_0^2} \right), \tag{1.47}$$

where ω_0 is the initial spin frequency and ω and $\dot{\omega}$ are the current values of spin frequency and spin frequency derivative. If we assume that $\omega_0 \gg \omega$, which means that the neutron star was rotating much faster at birth than at the current time, Equation (1.47) becomes:

$$\tau_{\mathrm{c}} = -\frac{\omega}{2\dot{\omega}} = \frac{P}{2\dot{P}} \simeq 1.6 \times 10^7\,\mathrm{yr} \left(\frac{P}{1\,\mathrm{s}} \right) \left(\frac{\dot{P}}{10^{-15}\,\mathrm{s\,s^{-1}}} \right)^{-1}. \tag{1.48}$$

The characteristic age in general represents an upper limit on the true neutron star age. Indeed for very young pulsars the assumption that $\omega_0 \gg \omega$ could not hold as their spin period can be still very close to the value at birth. Furthermore, for older pulsars, if one takes into account that the magnetic field decays over time the characteristic age could exceed the real age of the neutron star by orders of magnitude. In these cases the characteristic age fails to give a good estimate for the actual age.

1.9 Pulsar radio emission

As we have described in Section 1.7, the magnetosphere surrounding a neutron star is divided into a closed field line and an open field line region. According to the standard theory (Goldreich and Julian, 1969; Sturrock, 1971; Ruderman and Sutherland, 1975), in the open field line region particles can freely stream along the field lines. If charges are not continuously extracted from the surface the outflow of particles eventually generates vacuum or low-density regions, so-called plasma gaps, above the magnetic poles. Inside these gaps the force-free condition $\boldsymbol{E} \cdot \boldsymbol{B} = 0$ can be violated as the electric field parallel to the magnetic field lines is not perfectly screened anymore. New particles can therefore be extracted from the stellar surface and accelerated up to relativistic speed leading to an avalanche of pair production. The mechanism producing the radio emission in pulsars is not well understood, but it is believed that in this process bunches of $e^+\,e^-$ pairs produced in the gap oscillate coherently to produce radio emission at a certain distance from the polar cap see Section 1.9.3.

1.9.1 The polar cap model and radio emission geometry

The common picture is that this radio emission is generated in conical beams delimited by the open field lines around the magnetic poles. We can estimate the size of these conical beams by considering the size of the polar caps where the open field lines are anchored to the neutron star. The polar cap size depends on the amount
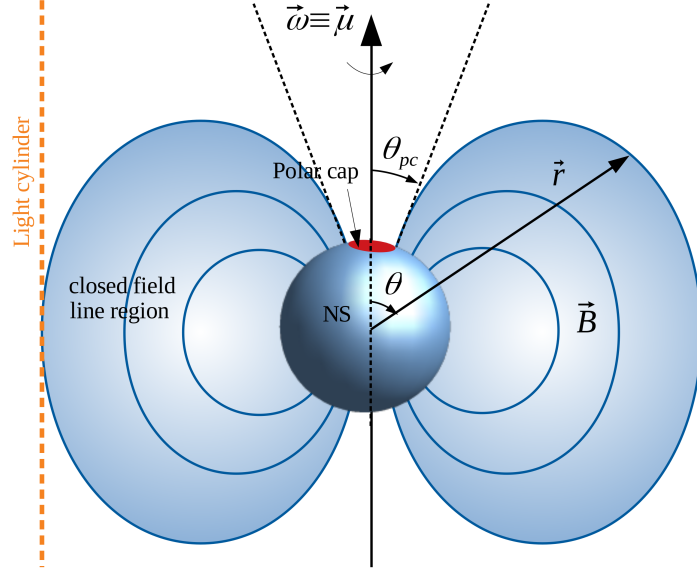
FIGURE 1.5: Sketch showing the geometry of the polar cap emission model for an aligned rotator. The polar cap region is delimited by the last closed field lines extending to the light cylinder. A magnetic field line can be parametrised as curve described by the coordinates r and θ (see text for more details).

of field lines that are opened which in turn depends on the size of the light cylinder (see Equation (1.26), i.e., on the spin period of the neutron star. To see this consider Figure 1.5 showing a sketch of a neutron star with the magnetic dipole moment aligned to the rotation axis of the star. We parametrise a magnetic field line as a curve $r = r(\theta)\hat{r}$ in space, where θ is the angle between the rotation axis and the vector r whose origin is located in the star's centre and pointing to a random position on the field line. The magnetic field vector B is therefore proportional to the tangent vector of the curve given by $d r/d\theta = (dr/d\theta)\hat{r} + r\hat{\theta}$, where $\hat{r}$ and $\hat{\theta}$ are the unit vectors in spherical coordinates. By considering the expression of the magnetic dipole in spherical coordinates (Equation (1.21)) we obtain the following equations:

$$\frac{\mathrm{d}r}{\mathrm{d}\theta} = \lambda(\theta)\frac{\mu}{r^3}2\cos(\theta), \tag{1.49}$$

$$r = \lambda(\theta)\frac{\mu}{r^3}\sin(\theta), \tag{1.50}$$

where $\lambda(\theta)$ represents a generic proportionality coefficient. If we consider the ratio between these two equations we end up with the following differential equation:

$$\frac{\mathrm{d}r}{r\mathrm{d}\theta} = \frac{2\cos\theta}{\sin\theta}. \tag{1.51}$$

We consider the last closed field line that delimits the polar cap region and is tangent to the light cylinder. The equation above can be integrated from a generic radius r corresponding to a generic θ to the light cylinder radius $r = r_{\mathrm{LC}}$ corresponding to $\theta = \pi/2$:

$$\int_r^{r_{\mathrm{LC}}} \frac{\mathrm{d}r'}{r'} = \int_\theta^{\pi/2} \frac{2\cos\theta'}{\sin\theta'}d\theta' \implies \ln\frac{r_{\mathrm{LC}}}{r} = 2\ln\frac{1}{\sin\theta}. \tag{1.52}$$

This implies:

$$\frac{\sin^2 \theta}{r} = \frac{1}{r_{\text{LC}}}. \tag{1.53}$$

Generally the quantity $\sin^2 \theta / r$ is constant along any dipolar field line. With this relation, by setting $r = R_{\text{NS}}$ we can find the angular radius of the polar cap on the neutron star surface:

$$\sin \theta_{\text{pc}} = \left(\frac{R_{\text{NS}}}{r_{\text{LC}}}\right)^{1/2} \tag{1.54}$$

$$\implies \theta_{\text{pc}} \sim \left(\frac{2\pi R_{\text{NS}}}{cP}\right)^{\frac{1}{2}} = 0.015 \, \text{rad} \, [0.87 \, \text{deg}] \left(\frac{R_{\text{NS}}}{11 \, \text{km}}\right) \left(\frac{P}{1 \, \text{s}}\right)^{-1/2}, \tag{1.55}$$

where we assumed that the the polar cap angular aperture is small i.e., $\theta_{\text{pc}} \lesssim 30°$ (Lorimer and Kramer, 2012).

We now consider an emission point on a field line at the edge of the polar cap, with coordinates $(r_{\text{em}}, \theta_{\text{em}})$. By using Equation (1.53) we can write:

$$\sin \theta_{\text{em}} = \left(\frac{r_{\text{em}}}{r_{\text{LC}}}\right)^{1/2} \implies \theta_{\text{em}} \sim \left(\frac{2\pi r_{\text{em}}}{cP}\right)^{1/2}, \tag{1.56}$$

where we have assumed again that the polar cap region is narrow and the emission point is close to the magnetic axis, i.e., $\theta_{\text{em}} \lesssim 30°$. As the field lines are curved, the opening angle ρ_{em} of the emission cone tangent to the field lines at the polar coordinates $(r_{\text{em}}, \theta_{\text{em}})$ is generally a little bigger than the polar angle θ_{em}. The relation between these two angles was found by Gangadhara and Gupta (2001), and for small angles reduces to:

$$\rho_{\text{em}} \sim \frac{3}{2}\theta_{\text{em}} \sim \left(\frac{9\pi r_{\text{em}}}{2cP}\right)^{\frac{1}{2}} = 0.069 \, \text{rad} \, [3.95 \, \text{deg}] \left(\frac{r_{\text{em}}}{100 \, \text{km}}\right)^{1/2} \left(\frac{P}{1 \, \text{s}}\right)^{-1/2}. \tag{1.57}$$

This gives the half-opening angle of the radio emission cone at the two magnetic polar regions of the neutron star. The modifications arising by taking into account a misaligned magnetic dipole moment and general relativistic effects are small and in general can be neglected (Kapoor and Shukre, 1998).

1.9.2 Pulse width

As the neutron star rotates, a pulse of radiation can be detected if our line of sight cuts through at least one of the radio beams. Therefore the features of the observed pulses depend on the geometry of the radio beam and the relative position at which the line of sight intercepts it. Having computed the angular aperture of the radio beam ρ_{em} (1.57), we can now set up a relation that links it to the intrinsic pulse half-width φ (expressed in radians) that an observer sees. We provide Figure 1.6, valid for a generic inclined rotator with χ the inclination angle between the magnetic dipole moment and the rotation axis, as a reference. Let's assume that as the star rotates our the line of sight intercepts the beam at an angular distance β from its centre. By assuming a sphere of unit radius, and considering the spherical triangle with sides $\chi + \beta$, χ and ρ_{em} we use the haversine law to write down the following relation:

$$\text{hav}\,(\rho_{\text{em}}) = \text{hav}\,(\chi + \beta - \chi) + \sin\,(\chi + \beta)\sin\,(\chi)\,\text{hav}\,(\varphi), \tag{1.58}$$
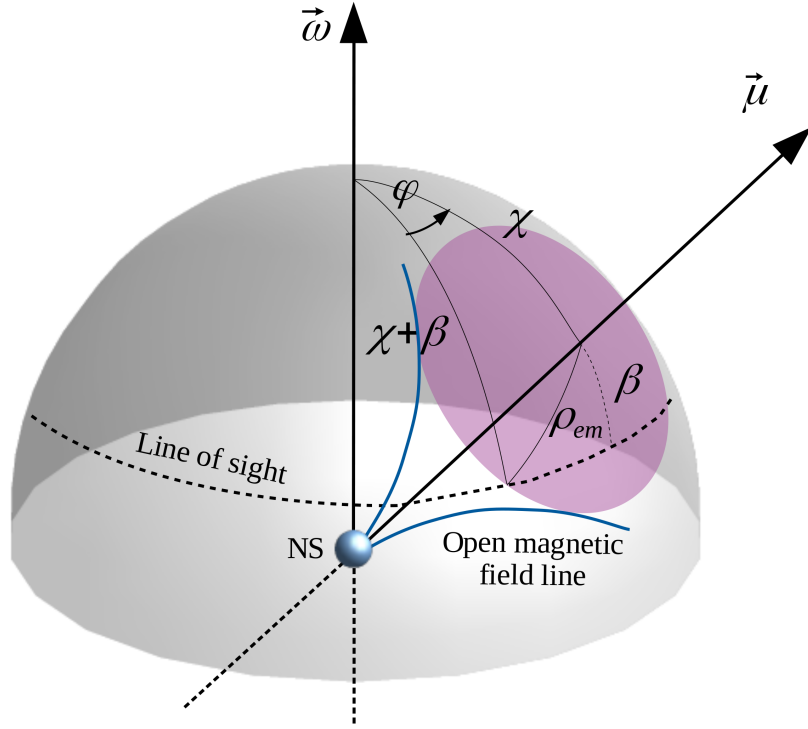
FIGURE 1.6: Radio beam geometry. The magnetic dipole moment μ is misaligned by an angle χ with respect to the rotation axis defined by ω. The line of sight cuts the emission cone at an angle β relative to the centre of the beam. The angular pulse width is given by φ.

where in general given an angle α the haversine function is defined as $\text{hav}(\alpha) = \sin^2(\alpha/2)$. Solving for the square sine of the half pulse width, we can write the following formula (see Maciesiak, Gil, and Ribeiro, 2011):

$$\sin^2\left(\frac{\varphi}{2}\right) = \frac{\sin^2\left(\frac{\rho_{em}}{2}\right) - \sin^2\left(\frac{\beta}{2}\right)}{\sin(\chi + \beta)\sin(\chi)} \tag{1.59}$$

$$\implies \varphi = 2\arcsin\sqrt{\frac{\sin^2\left(\frac{\rho_{em}}{2}\right) - \sin^2\left(\frac{\beta}{2}\right)}{\sin(\chi + \beta)\sin(\chi)}}. \tag{1.60}$$

To find the pulse width in time units we have to re-normalise for the spin period:

$$w = 2\varphi\frac{P}{2\pi} = 2\frac{P}{\pi}\arcsin\sqrt{\frac{\sin^2\left(\frac{\rho_{em}}{2}\right) - \sin^2\left(\frac{\beta}{2}\right)}{\sin(\chi + \beta)\sin(\chi)}}. \tag{1.61}$$

A simplified formula for the pulse width can be found if one assumes that the line of sight intercepts the radio beam in its centre (i.e. with $\beta = 0$) and that ρ_{em} and φ are small ($\lesssim 30°$). In this case we can approximate the relation in Equation (1.59) as:

$$\left(\frac{\varphi}{2}\right)^2 \sim \frac{\left(\frac{\rho_{em}}{2}\right)^2}{\sin^2(\chi)} \implies \varphi \sim \frac{\rho_{em}}{\sin(\chi)}, \tag{1.62}$$
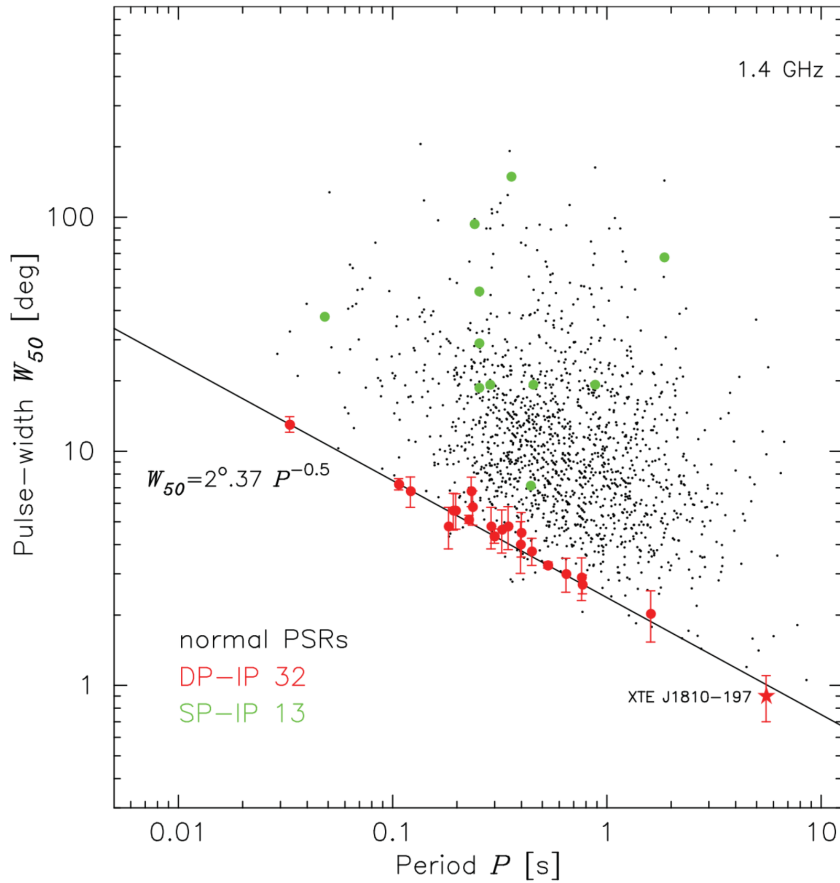
FIGURE 1.7: Plot of pulse-widths w_{50} measured at the 50% of the maximum pulse amplitude versus the spin period P for 1450 normal pulsars detected at 1.4 GHz (excluding millisecond and other recycled pulsars) from the ATNF database (Manchester et al., 2005). The figure shows also 21 cases of pulsars being almost orthogonal rotators and manifesting inter-pulses due to our line of sight crossing both magnetic poles (red dots) and 11 cases of almost aligned pulsars where only one magnetic pole is visible but manifesting interpulses (green dots). The black solid line represents the fit of the orthogonal-rotator pulsars (red dots) (see Maciesiak and Gil, 2011; Maciesiak, Gil, and Melikidze, 2012, for more details) (Figure taken from Maciesiak, Gil, and Melikidze, 2012).

which leads to:

$$w = 2\varphi \frac{P}{2\pi} \sim 2 \frac{\rho_{\rm em}}{\sin \chi} \frac{P}{2\pi} \sim \frac{2}{\sin \chi} \left(\frac{r_{\rm em} P}{2\pi c} \right)^{\frac{1}{2}} \tag{1.63}$$

$$\simeq \frac{0.015 \, {\rm s}}{\sin \chi} \left(\frac{r_{\rm em}}{100 \, {\rm km}} \right)^{1/2} \left(\frac{P}{1 \, {\rm s}} \right)^{1/2}. \tag{1.64}$$

Figure 1.7 shows the relation between the pulse-widths w_{50} measured at the 50% of the maximum pulse amplitude versus the spin period P for 1450 normal pulsars detected at 1.4 GHz (excluding millisecond pulsars). The black line represents a fit of the pulsars showing interpulses and predicted to be orthogonal rotators, i.e. with $\chi \sim 90°$. Note that in Equation 1.63, the $\sin \chi$ term corrects for the inclination angle. For a given angular aperture, $\rho_{\rm em}$, of the emission cone, as the inclination angle decreases, the observed pulse width increases. This explain why in Figure 1.7 the pulsars that are likely orthogonal rotators (red dots) are found at the extreme bottom part of the diagram and more aligned rotators are scattered above depending on the inclination angle. For example, if we observe an aligned rotator ($\chi \sim 0°$) we do not see pulsed emission since our line of sight is always inside the radio beam. In this case, according to the formula above the pulse width goes to ∞ but this simply indicates that we are seeing continuous emission from the polar cap. Furthermore, as $-\rho_{\rm em} < \beta < \rho_{\rm em}$, the angle $\chi + \beta$ could become negative. This occur only when $\rho_{\rm em} > \chi$, and in that case it happens that the line of sight is always inside the beam and therefore also in this case no pulsation is observed. Some periodic modulation in the observed pulse profile can be still observed if the intensity profile of the beam is not uniform.

To obtain an estimate of the emission cone aperture $\rho_{\rm em}$ from the pulse width w measurements one can in principle use Equation (1.61). However, deriving an estimate of the intrinsic pulse width w is not always trivial. The observed pulse width usually differ from the intrinsic pulse width due to pulse broadening as the radio waves propagates through the interstellar medium (see Section 1.10). Furthermore Pulsars exhibit different pulse profiles and widths due to various intrinsic properties, including magnetospheric dynamics, emission processes, and rotational irregularities. These differences make it challenging to predict or measure the pulse width accurately, especially for pulsars with complex emission patterns (see Lorimer and Kramer, 2012). Moreover, we generally lack information on the relative angular position between the rotation axis, magnetic axis and line of sight. For more details see also Chapter 3 in Lorimer and Kramer (2012).

Overall even if this theoretical picture does not take into account all the complexity that might arise in reality, it predicts a relation $\rho_{\rm em} \sim P^{-1/2}$ which has been observed for samples of pulsars with reliable measurements of the pulse width (see Figure 1.7 and Rankin, 1993; Maciesiak and Gil, 2011; Maciesiak, Gil, and Melikidze, 2012). By using a sample of pulsars observed by the Parkes Murriyang radio telescope Johnston and Karastergiou (2019) found a shallower relation $\rho_{\rm em} \sim P^{-0.3}$. However the scatter of data in the pulse-width – spin-period plane is large and is partly produced by the unknown magnetic inclination angle χ as can be seen from Equation (1.63). As χ is unknown, its effect is difficult to take into account and affects the correlation between these quantities. This work also suggests that a narrow range of emission heights $r_{\rm em}$ between around 200 km and 400 km seems to be preferred.
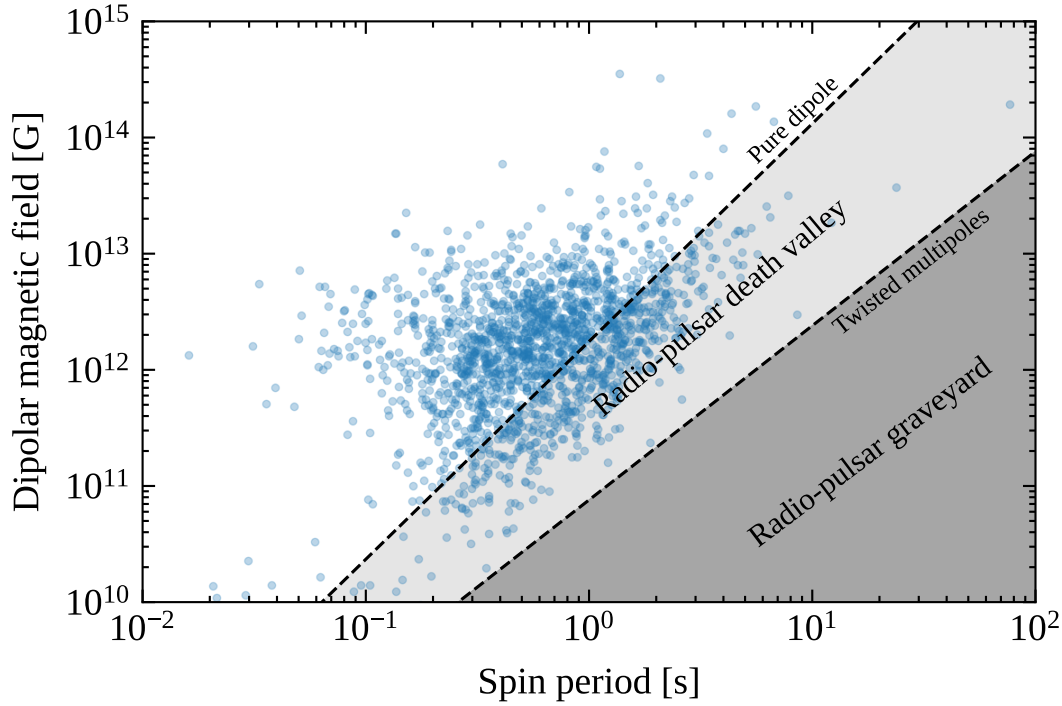
FIGURE 1.8: Radio emission death valley delimited by the two extreme death lines (*dashed lines*) defined in Equations (1.71) for a pure dipole configuration, and (1.72) for a twisted multipole configuration, respectively. The dots represent the observed radio pulsars from the ATNF catalog (Manchester et al., 2005), where the dipolar magnetic field values are estimated from the timing properties (P and $\dot{P}$) using Equation 1.46.

1.9.3 Radio emission mechanism and radio death valley

If a gap devoid of plasma is formed above the polar cap region an electric potential difference is generated along a magnetic field line. Following Ruderman and Sutherland (1975) we can consider a narrow gap above the polar region extending from the neutron star surface up to a height H. A given potential difference ΔV will accelerate an electron (or a positron) up to an energy $\gamma m_e c^2 = e\Delta V$ which implies a Lorentz factor of:

$$\gamma = \frac{e\Delta V}{m_e c^2}. \tag{1.65}$$

By moving along the curved field lines the electron will emit curvature photons of energy (Jackson, 1998):

$$hf = \frac{3}{4\pi} \frac{hc}{r_c} \gamma^3 = \frac{3}{4\pi} \frac{hc}{r_c} \left(\frac{e\Delta V}{m_e c^2} \right)^3, \tag{1.66}$$

where h is the Planck constant, f is the photon frequency and r_c denotes the curvature radius of the magnetic field line. The condition that such a curvature photon with energy $hf > 2m_e c^2$ moving through a magnetic field produces an electron-positron

pair inside the gap is given by (see Ruderman and Sutherland, 1975; Chen and Ruderman, 1993):

$$\frac{hf}{2m_ec^2}\frac{B_\perp}{B_{\text{QED}}} \sim \frac{1}{15}. \tag{1.67}$$

Here $B_{\text{QED}} = 2\pi m_e^2 c^3/(eh)$ is the Schwinger magnetic field, defined as the magnetic field strength for which a cyclotron photon emitted by an electron gyrating around a field line will have an energy equal to m_ec^2. $B_\perp$ is the component of the magnetic field perpendicular to the direction of propagation of the photon. Due to relativistic beaming the curvature photons are emitted along the direction of motion of the electron, i.e., almost parallel to the magnetic field lines, within an angle $\sim 1/\gamma$. Suppose that a curvature photon is emitted close to the stellar surface, where the magnetic field strength is B_s. As the photon propagate inside the gap the component $B_\perp$ will grow up to an approximate maximum value given by $B_\perp \sim \frac{H}{r_c}B_s$. With all these considerations we can substitute Equation (1.66) into Equation (1.67) to find the relation:

$$\frac{3}{8\pi}\left(\frac{e\Delta V}{m_ec^2}\right)^3 \frac{h}{m_ec}\frac{H}{r_c^2}\frac{B_s}{B_{\text{QED}}} \sim \frac{1}{15}. \tag{1.68}$$

This equation defines the minimum potential difference that is required to trigger a pair cascade inside the gap. According to the standard theory if a neutron star is not able to generate such electric potential difference in its magnetosphere it cannot be visible as a radio pulsar. On the other hand, triggering the pair cascade will also limit any further growth of ΔV above the value established by Equation (1.68) as the produced plasma will progressively screen the electric field.

Assuming an aligned dipolar configuration for the magnetic field, due to the unipolar induction mechanism (see Section 1.7), the maximum potential difference that can be generated along an open field line in the magnetosphere is given by the difference between the electric potential on the stellar surface given by Equation (1.25)) and the electric potential of the interstellar medium to which the open field line is connected at infinity. The interstellar electric potential can be assumed to vanish. From Equation (1.25), the maximum potential difference can be established on the last open field line at the edge of the polar cap, i.e. at a coordinate $\theta = \theta_{\text{pc}}$. By assuming $V_P = 0$ we have:

$$\Delta V_{\text{max}} \simeq \frac{\omega\mu\sin^2\theta_{\text{pc}}}{2R_{\text{NS}}c} = \frac{\omega B_P R_{\text{NS}}^2 \sin^2\theta_{\text{pc}}}{2c} = \frac{\omega^2 B_P R_{\text{NS}}^3}{2c^2}. \tag{1.69}$$

where we used Equation (1.44) and used the relation (1.54). This maximum potential difference depends solely on the dipole component of the field as it is the one determining the open field line configuration at large distance from the stellar surface, and is valid even if multipoles dominate near the stellar surface (Ruderman and Sutherland, 1975). If the maximum potential difference ΔV_{max} that can be produced in the magnetosphere is lower than the threshold defined by Equation (1.68) the pair cascade can not be triggered and the neutron star is expected to be radio quiet. Therefore a necessary condition that has to be satisfied to turn on the radio emission is the following:

$$\frac{3}{8\pi}\left(\frac{e\Delta V_{\text{max}}}{m_ec^2}\right)^3 \frac{h}{m_ec}\frac{H}{r_c^2}\frac{B_s}{B_{\text{QED}}} \gtrsim \frac{1}{15} \tag{1.70}$$

Depending on the assumptions that one makes on the magnetic field configuration near the surface and the gap geometry, that is the values of B_s, r_c and H, the condition (1.70) defines a threshold for the turning on and off of the radio emission in a pulsar. As the suppression of the radio emission depends on the particular magnetospheric properties of each pulsar, relation (1.70) defines a "death valley" for the radio pulsars (see Chen and Ruderman, 1993). We can consider two extreme cases. For an ideal case of a pure dipole configuration we consider $B_s = B_p$, the curvature radius for an open field line can be approximated as $r_c \sim r_{LC} \sin \theta_{pc} = (R_{NS} r_{LC})^{1/2}$ and the gap height can be taken to be equal to the polar cap radius $H \sim R_{NS} \sin \theta_{pc} = R_{NS} (R_{NS}/r_{LC})^{1/2}$ (Chen and Ruderman, 1993). With these order of magnitude assumptions the following condition for the radio emission turning on is derived from Equation (1.70):

$$B_p \gtrsim 2.0 \times 10^{12}\,\text{G} \left(\frac{R_{NS}}{10^6\,\text{cm}} \right)^{-19/8} \left(\frac{P}{1\,\text{s}} \right)^{15/8}. \tag{1.71}$$

If instead multipoles dominate the magnetic field configuration near the surface, the value B_s could exceed the one of a pure dipole. We can introduce the parametrisation $B_s = b B_p$ where b is a numerical factor that takes into account deviations from the dipolar field value. Furthermore the curvature radius r_c in general scales inversely with the multiple order. We can consider an extreme case where the magnetic field near the surface is very twisted for the presence of strong multipole components. In this case for example we can assume $r_c \sim H \sim R_{NS}$. In this case, the condition for the radio emission becomes:

$$B_p \gtrsim 8.3 \times 10^{10}\,\text{G}\, b^{-1/4} \left(\frac{R_{NS}}{10^6\,\text{cm}} \right)^{-2} \left(\frac{P}{1\,\text{s}} \right)^{3/2}. \tag{1.72}$$

Given a neutron star with a spin period P, Equation (1.71) and (1.72) give the minimum polar magnetic field strength that the star should have to trigger the radio emission. Neutron stars with a lower magnetic field are considered radio-dead (see bottom right corner in Figure 1.8).

1.9.4 Energetics

As we have seen in Section 1.7 and Section 1.8, a neutron star behaves as a unipolar inductor and spins down due to magnetic dipole breaking. It converts its rotational kinetic energy into electrical potential energy that ultimately leads to particle acceleration and electromagnetic emission. The rate of loss of rotational energy is given by Equation (1.27). This represents the total budget of rotational power that can be converted into electromagnetic power by the unipolar inductor process. As the potential difference that is established inside the magnetosphere is created at the expense of the rotational energy we expect a relation between the maximum potential difference and the rotational energy loss rate. By expressing the polar magnetic field strength in Equation (1.69) as a function of the timing properties P and $\dot{P}$ through Equation (1.46) the maximum potential difference inside the magnetosphere can be expressed as:

$$\Delta V_{max} \simeq \left(\frac{3}{2c} \frac{dE_{rot}}{dt} \right)^{1/2}. \tag{1.73}$$

We neglect the effect of the inclination angle which gives a correction of a factor 2 at most. The energy of the particles accelerated by this potential is given by $e\Delta V_{\mathrm{max}}$. This is the amount of energy that each particle can radiate away in the form of electromagnetic radiation. Therefore the rotational energy is first converted into electrical potential energy which leads to particle acceleration and ultimately is radiated away in the form of electromagnetic radiation. From this arguments it is reasonable to assume a link between the radio luminosity and the spin-down power. We consider a generic power-law dependence of the intrinsic radio luminosity L of a pulsar on the rotational energy loss rate as:

$$L = L_0 \left(\frac{\dot{E}_{\mathrm{rot}}}{\dot{E}_{\mathrm{rot},0}} \right)^{\epsilon} = \mathcal{L}_0 \left(\frac{\dot{P}}{P^3} \right)^{\epsilon}, \tag{1.74}$$

where L_0 is the luminosity at the reference rotational power $\dot{E}_{\mathrm{rot},0}$ and:

$$\mathcal{L}_0 = L_0 \left(\frac{4\pi^2 I_{\mathrm{NS}}}{\dot{E}_{\mathrm{rot},0}} \right)^{\epsilon}. \tag{1.75}$$

Considering a typical pulsar radio luminosity $L_0 \sim 10^{29}$ erg s^{-1} corresponding to $\dot{E}_{\mathrm{rot},0} = 10^{30}$ erg s^{-1} (see Szary et al., 2014), $\epsilon = 0.5$ and default values for the moment of inertia of a neutron star $I_{\mathrm{NS}} \sim 10^{45}$ g cm^2, we obtain $\mathcal{L}_0 \sim 3 \times 10^{35}$. As the pulsar spins down it converts part of the rotational power in the form of radio waves. As the radio luminosity should not exceed the spin-down energy loss rate we can also express the equation above in terms of an efficiency η:

$$L = \eta(\dot{E}_{\mathrm{rot}})\dot{E}_{\mathrm{rot}}. \tag{1.76}$$

The efficiency is a function of the rotational power itself and in principle should not exceed 1. Equation (1.74) represents the intrinsic total radio luminosity emitted by the two radiation beams. This intrinsic luminosity is however difficult to derive from observations and instead, a pseudo-luminosity is usually estimated as we will discuss in Section 1.11.

1.10   Pulse propagation in the interstellar medium

The intensity, shape and duration of a detected radio pulse are influenced by a variety of factors which depend on both the propagation in the interstellar medium and instrumental effects. In the following I will discuss these aspects individually.

Suppose that a pulsar with spin period P and emitting two beams with an intrinsic radio luminosity L is located at a distance d from Earth. If the radio signal from this pulsar were to travel in complete vacuum, the pulse of radiation would have a width w given by Equation (1.59) and the flux of radiation reaching the Earth would be:

$$S = \frac{L}{\Omega_{\mathrm{em}} d^2}. \tag{1.77}$$
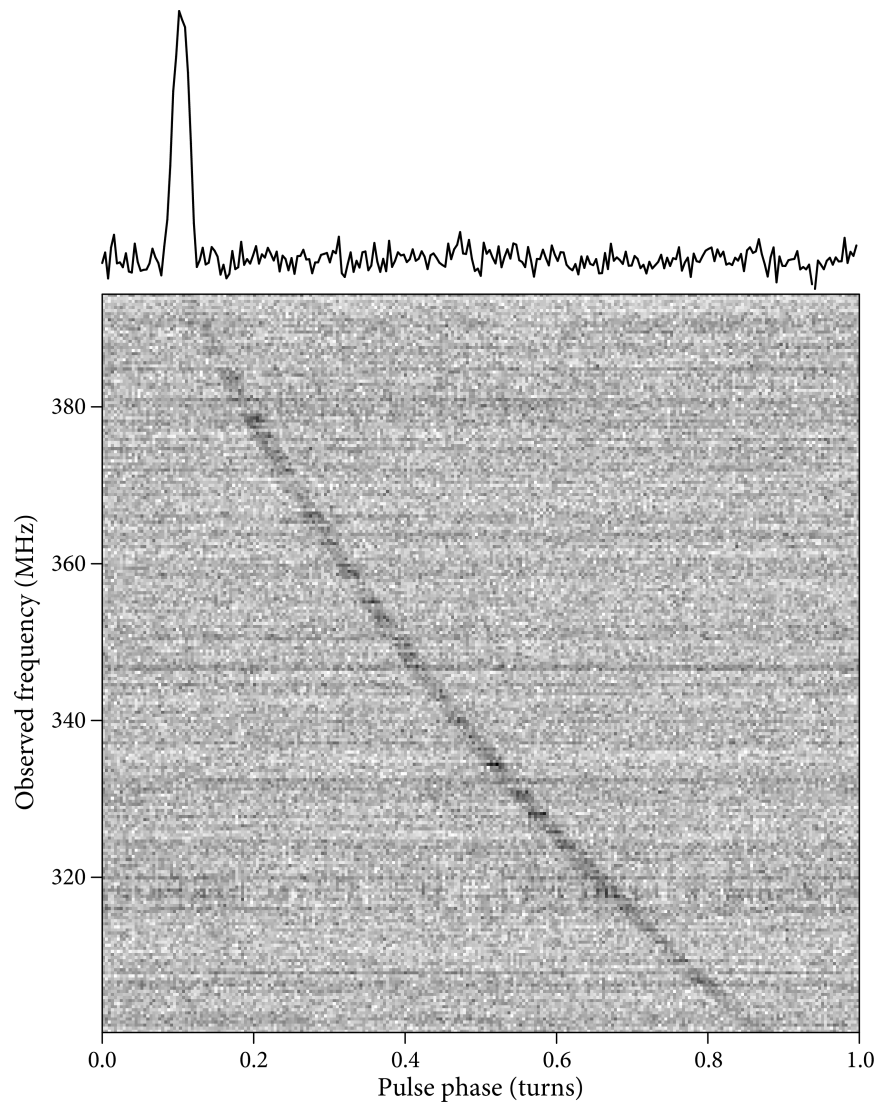
FIGURE 1.9: Pulse dispersion of the pulsar J1800+5034. The dark band in the frequency versus phase diagram shows the delay in arrival time of radio waves of different frequency. The "de-dispersed" band-integrated pulse profile is shown at the top. (Figure taken from https://www.cv.nrao.edu/~sransom/web/Ch6.html).
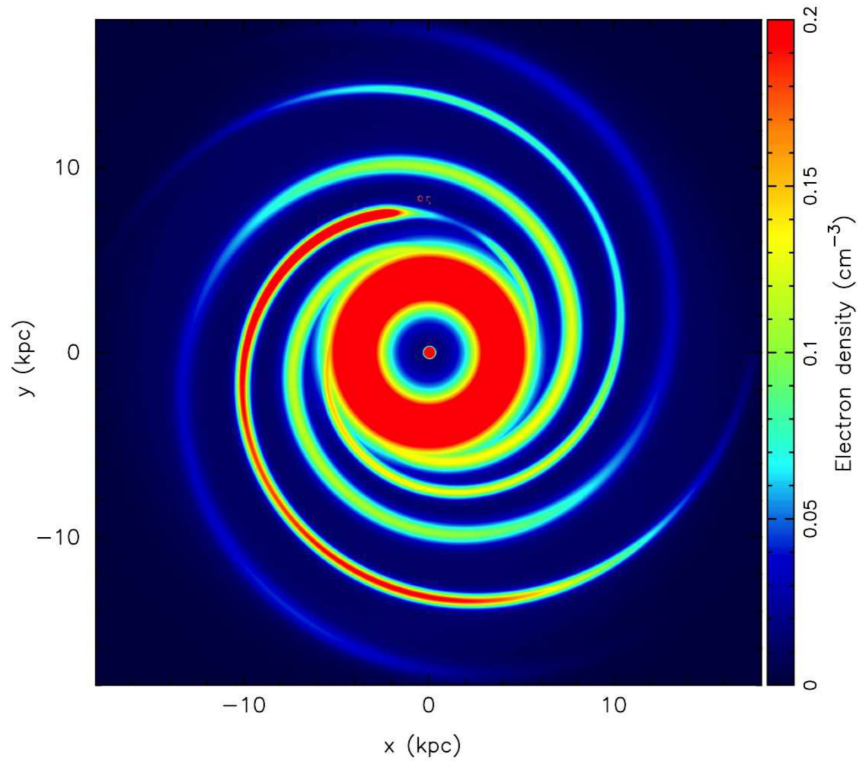
FIGURE 1.10: Electron density in the Galactic plane ($z = 0$) for the YMW16 model (Yao, Manchester, and Wang, 2017). The Galactic Centre is at the origin and the Sun is at $x = 0$, $y = 8.5$ kpc. The dense annulus of central radius 4 kpc represents the Galactic molecular ring. Spiral arms have a pre-determined logarithmic spiral form and generally decay exponentially at large Galactocentric radii. Exceptions are in the Carina and Sagittarius arms where there are over-dense and under-dense regions in Carina and Sagittarius respectively. The Gum Nebula and Local Bubble features are faintly visible near the position of the Sun (see Yao, Manchester, and Wang, 2017, for more details) (Figure taken from Yao, Manchester, and Wang, 2017).

Here Ω_{em} is the total solid angle covered by the two beams which are given by (see Lorimer and Kramer, 2012):

$$\begin{aligned}
\Omega_{\text{em}} &= 2 \times 2\pi \int_0^{\rho_{\text{em}}} \sin\rho\,\mathrm{d}\rho \\
&= 4\pi\left(1 - \cos\rho_{\text{em}}\right) \\
&= 8\pi \sin^2\left(\frac{\rho_{\text{em}}}{2}\right).
\end{aligned} \tag{1.78}$$

The factor 2 comes from the two radio beams corresponding to each magnetic pole. In the last step we used the trigonometric identity $(1 - \cos\rho_{\text{em}}) = 2\sin^2(\rho_{\text{em}}/2)$.

Radio telescopes are sensitive to a limited range of frequencies around a central frequency f_0. The radio flux at a given frequency can be evaluated by considering the spectral shape of the signal. Radio pulsars are usually characterised by a power-law spectral shape (see Jankowski et al., 2018). Therefore to compute the flux density at a given frequency f we can use the following relation:

$$S_f(f) = S_{f,0}\left(\frac{f}{f_0}\right)^\alpha = \frac{S(\alpha+1)}{f_{\text{max}}^{\alpha+1} - f_{\text{min}}^{\alpha+1}} f^\alpha, \tag{1.79}$$

where we used the fact that S is the flux integrated over the entire frequency range between f_{min} and f_{max}:

$$S = \int_{f_{\text{min}}}^{f_{\text{max}}} S_f(f)\,\mathrm{d}f = S_{f,0}\frac{f_0^{-\alpha}}{\alpha+1}\left(f_{\text{max}}^{\alpha+1} - f_{\text{min}}^{\alpha+1}\right). \tag{1.80}$$

We can also define the fluence of this signal by integrating the flux over the interval of time defined by the spin period:

$$\mathcal{F}_f(f) = \int_0^P S_f(f,t)\,\mathrm{d}t = S_f(f)w, \tag{1.81}$$

where for simplicity we considered a rectangular top-hat pulse shape with intrinsic width w. The fluence has units of [erg cm^{-2}] and is a measure of the total energy carried by a radiation pulse passing through a unit surface area.

However, to reach the Earth the signal has to propagate through the interstellar medium which is characterised by clumpy regions with different densities where gas can be completely ionised. In the presence of ionised gas, the electromagnetic interaction between the photons and the free electrons causes a delay in the propagation of the light, which is a function of photon frequency. More energetic photons tend to push past the free electrons with little effect on their propagation speed, whereas lower frequency photons (like radio waves) are more significantly delayed. We quantify these effects by introducing the refraction index n, which for a non-magnetised, non-relativistic and ionised plasma is given by:

$$n(f) = \sqrt{1 - \left(\frac{f_{\text{p}}}{f}\right)^2}, \tag{1.82}$$

where f_{p} is the so-called plasma frequency for a non-magnetised plasma:

$$f_{\text{p}} = \left(\frac{e^2 n_e}{\pi m_e}\right)^{1/2} \simeq 8.98\,\text{kHz}\left(\frac{n_e}{1\,\text{cm}^{-3}}\right)^{1/2}. \tag{1.83}$$

The delay in the arrival time of a wave of frequency f with respect to the arrival time in vacuum can be estimated to be (see Lorimer and Kramer, 2012):

$$
\begin{aligned}
\Delta t(f) &= \left(\int_0^d \frac{\mathrm{d}l}{n(f)c} \right) - \frac{d}{c} \simeq \int_0^d \frac{\mathrm{d}l}{c} \frac{f_{\mathrm{p}}^2}{2f^2} \\
&= \frac{e^2}{2\pi m_e c} f^{-2} \int_0^d \mathrm{d}l n_e(l) \\
&\simeq 0.41\,\mathrm{s} \left(\frac{f}{1\,\mathrm{GHz}} \right)^{-2} \frac{DM}{100\,\mathrm{pc\,cm^{-3}}},
\end{aligned}
\tag{1.84}
$$

where we assumed $f_{\mathrm{p}} \ll f$. The integral

$$
DM = \int_0^d \mathrm{d}l n_e(l),
\tag{1.85}
$$

is performed along the line of sight l, over a distance d and computes the electron column density along the line of sight which defines the dispersion measure DM.

To recover the original pulse shape, the detected signal has to be de-dispersed by re-aligning the observations of delayed pulses at different frequencies (see Figure 1.9). However, a radio detector has a finite frequency resolution. Usually, the total bandwidth of the instrument is divided into different channels, each one sensitive to photons in a particular range of frequencies Δf_{ch}. Since photons collected by a given spectral channel cannot be distinguished in frequency, the de-dispersion cannot be applied per channel. To obtain an estimate of the residual intra–channel dispersion smearing contribution to the observed pulse width we instead differentiate Equation (1.84) with respect to the frequency f to find the amount of delay per unit frequency interval and multiply by the minimum frequency bin Δf_{ch} of the instrument:

$$
\begin{aligned}
\tau_{DM} &= \left| \frac{\mathrm{d}\delta t}{\mathrm{d}f} \right| \Delta f_{\mathrm{ch}} \\
&= \frac{e^2}{\pi m_e c} f^{-3} \Delta f_{\mathrm{ch}} DM \\
&\simeq 0.83\,\mathrm{ms} \left(\frac{f}{1\,\mathrm{GHz}} \right)^{-3} \frac{\Delta f_{\mathrm{ch}}}{1\,\mathrm{MHz}} \frac{DM}{100\,\mathrm{pc\,cm^{-3}}}.
\end{aligned}
\tag{1.86}
$$

Furthermore, the electron density in the medium is not homogeneous and shows variations over a wide range of length scales. The distribution of scales of these variations is affected by the presence of turbulence in the medium, which introduces a relative motion between the source, the inhomogeneous clumps and the observer. As the electromagnetic signal propagates, it encounters several such irregularities depending on the travelled distance d. At each encounter, it is delayed and deflected in a frequency-dependent way. The result is that a sharp pulse emitted by a point source is detected as a scattered-broadened pulse whose intensity decays exponentially with a characteristic time τ_{scat} which depends on the signal frequency and distance of the source (see Lorimer and Kramer, 2012). By assuming a Kolmogorov turbulence power spectrum for the inhomogeneity length scales (Kolmogorov, 1941) the broadening timescale due to scattering is given by:

$$
\tau_{\mathrm{scat}} \propto d^{2.2} f^{-4.4}.
\tag{1.87}
$$

Some studies show evidence of a relation between the scattering broadening timescale and the dispersion measure. For example, by using a sample of scattering measurements from 124 pulsars detected at 327 MHz by the Ooty Radio Telescope in Southern India, Krishnakumar et al. (2015) found the following empirical relation:

$$\tau_{\text{scat}} = 3.6 \times 10^{-9} \, \text{s} \left(\frac{DM}{1 \, \text{pc cm}^{-3}} \right)^{2.2} \left[1.0 + 1.94 \times 10^{-3} \left(\frac{DM}{1 \, \text{pc cm}^{-3}} \right)^{2.0} \right]. \quad (1.88)$$

This empirical relation is useful to estimate the scattering timescale for sources which only have a measurement of the dispersion measure.

Additionally, there are instrumental effects that contributes to the broadening of the pulse shape during the digital processing of the radio signal. The most relevant effect is given by the detector sampling time which limits the time resolution of the pulse and introduce an uncertainty in the pulse duration given by τ_{samp}.

As a result of the effects discussed here, the observed duration w_{obs} of a radio pulse is larger than the intrinsic duration w and is given by the quadratic sum of the contributions described above as reported in Cordes and McLaughlin (2003):

$$w_{\text{obs}} \simeq \sqrt{w^2 + \tau_{DM}^2 + \tau_{\text{sc}}^2 + \Delta t_{\text{samp}}^2}, \quad (1.89)$$

As the pulse propagates, gets dispersed and broadens, its fluence is defined by Equation (1.81) is conserved. Therefore the flux of the pulse detected on Earth can be found by:

$$S_f(f) w = S_{f,\text{obs}}(f) w_{\text{obs}} \implies S_{f,\text{obs}}(f) = \frac{S_f(f) w}{w_{\text{obs}}}. \quad (1.90)$$

1.11 Pulsar pseudo-luminosity

By inverting the steps described in the last section, one should in principle be able to estimate the intrinsic luminosity of a pulsar from the observed quantities. Usually in pulsar radio surveys the flux $S_{f,\text{mean}}$ averaged over a period and at the observation frequency f_0 is reported. Assuming for simplicity a rectangular pulse shape with peak flux $S_{f,\text{obs}}$ and observed width w_{obs} this is given by:

$$S_{f,\text{mean}}(f_0) = \frac{\int_0^P S_{f,\text{obs}}(f_0, t) \, \text{d}t}{P} = \frac{S_{f,\text{obs}}(f_0) w_{\text{obs}}}{P}. \quad (1.91)$$

To recover the intrinsic flux $S(f_0)$ we need to find the intrinsic pulse width. From Equation (1.89) we have:

$$w \simeq \sqrt{w_{\text{obs}}^2 - \tau_{DM}^2 - \tau_{\text{scat}}^2 - \tau_{\text{samp}}^2}. \quad (1.92)$$

By combining Equation (1.90) and Equation (1.91) with the previous equation we then find:

$$S_f(f_0) \simeq S_{f,\text{mean}}(f_0) \frac{P}{\sqrt{w_{\text{obs}}^2 - \tau_{DM}^2 - \tau_{\text{scat}}^2 - \tau_{\text{samp}}^2}} = \frac{S_{f,\text{mean}}(f_0)}{\delta}, \quad (1.93)$$

where here we define the intrinsic duty cycle $\delta \equiv w/P$. Finally using Equation (1.77) and Equation (1.78) we obtain for the intrinsic luminosity density:

$$L_f(f_0) \simeq \frac{8\pi}{\delta} \sin^2 \left(\frac{\rho_{\rm em}}{2} \right) d^2 S_{f,{\rm mean}}(f_0). \tag{1.94}$$

And integrating over the frequencies (Equation (1.80)) we obtain the bolometric intrinsic radio luminosity:

$$L \simeq \frac{8\pi}{\delta} \sin^2 \left(\frac{\rho_{\rm em}}{2} \right) d^2 S_{\rm mean}(f_0) \frac{f_0^{-\alpha}}{\alpha+1} \left(f_{\rm max}^{\alpha+1} - f_{\rm min}^{\alpha+1} \right). \tag{1.95}$$

However, a precise estimate of the intrinsic radio luminosity is not trivial for the following reasons:

- In most cases estimating the angular aperture $\rho_{\rm em}$ of the radio beam and the intrinsic pulse width w can be challenging (see also Section 1.9.2). Pulsars often exhibit intricate emission patterns, making it highly complicated to deduce the emission geometry and accurately determine the pulse width. One possibility to estimate the beam angular aperture from the spin period is to use the relation given by Equation (1.57), which has been confirmed by observation despite some uncertainties (Rankin, 1993; Maciesiak, Gil, and Melikidze, 2012). Additionally, removing the broadening effects described above from the observed pulse width $w_{\rm obs}$ is not a straightforward task, as it can be difficult to differentiate these effects from the inherent properties of the pulse shape for instrumental limitations in the time and frequency resolution (Lorimer and Kramer, 2012).

- Accurate modelling of radio pulsar spectra is not always possible. Radio spectra appear to be different for every pulsar. As shown in Jankowski et al. (2018) in many cases a single power-law model describes the observed spectrum well but many pulsars deviate from this simple behaviour.

- The estimate of the distance d is subject to many uncertainties. For radio pulsars, the main way to estimate the distance relies on models for the electron density distribution in the Galaxy. Three main models have been developed over the years from (see Taylor and Cordes, 1993; Cordes and Lazio, 2002; Yao, Manchester, and Wang, 2017). By knowing the DM value and the sky position of the pulsar, these models predict the amount of free electrons along the line of sight and give a rough estimate of the distance. However, due to the complexity of modelling the Galactic structures, these models suffer from a lot of uncertainties. The relative error on the distance estimated from the DM values could indeed be around 40 % (Yao, Manchester, and Wang, 2017).

Due to these uncertainties fiducial values are usually adopted for the unknown quantities (see Lorimer and Kramer, 2012). For example by considering a fixed beam aperture of $\rho_{\rm em} \sim 6\,{\rm deg}$, the ratio between the intrinsic pulse width and the spin period (the duty cycle) of $\delta = 0.04$, a radio frequency range from $f_{\rm min} = 10\,{\rm MHz}$ to $f_{\rm max} = 100\,{\rm GHz}$ (see for example Izvekova et al., 1981; Morris et al., 1997), and a spectral index $\alpha = -1.6$ (Jankowski et al., 2018), we have a pseudo luminosity of:

$$L_{f,{\rm pseudo}}(1.4\,{\rm GHz}) \simeq 7.4 \times 10^{27}\,{\rm erg\,s}^{-1} \left(\frac{S_{f,{\rm mean}}(1.4\,{\rm GHz})}{1\,{\rm mJy}} \right) \left(\frac{d}{1\,{\rm kpc}} \right)^2. \tag{1.96}$$
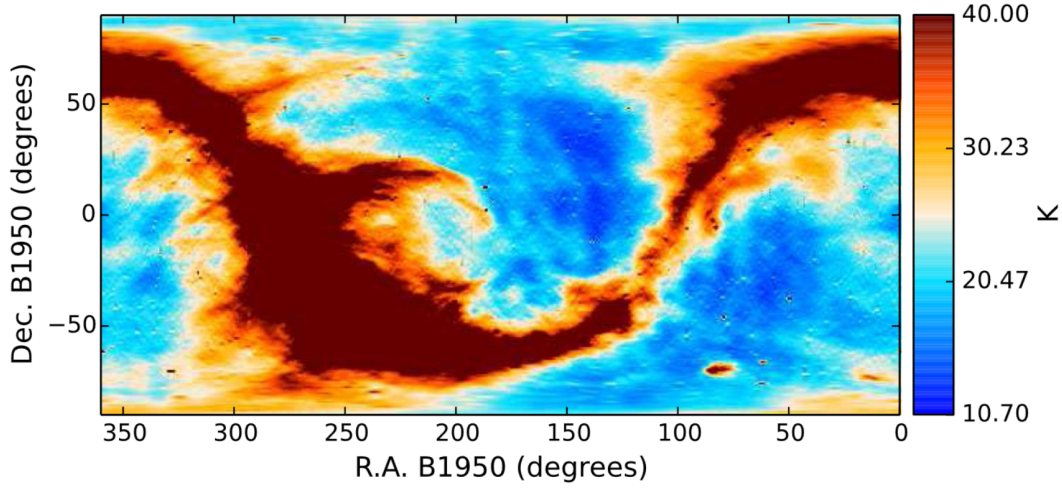
FIGURE 1.11: Sky temperature distribution at 408 MHz from Haslam et al. (1982), refined by Remazeilles et al. (2015) (Figure taken from Remazeilles et al., 2015).

However note that performing a population study and applying this expression to different pulsars without taking into account differences in their intrinsic duty cycles and beam geometries will mask eventual intrinsic correlations between the intrinsic luminosity and the spin-down properties like P, $\dot{P}$ and $\dot{E}_{rot}$. This is because ρ_{em} and w have a non-trivial dependence on the spin period (see Section 1.9). For example Posselt et al. (2023) recently showed that the radio pseudo-luminosity has a shallow dependence on the spin-down energy loss $L_{f,pseudo} \propto \dot{E}_{rot}^{0.15}$ (see also Szary et al., 2014). However, this might not be a proof of the absence of intrinsic correlations between the intrinsic radio luminosity and the rotational properties.

1.12 Pulsar detection, the radiometer equation

In this section we will describe how we can model the sensitivity of a radio telescope which is a crucial aspect to determine which pulsars are detectable. We follow Lorimer and Kramer (2012).

Consider the observation of a signal with peak flux density $S_{f,obs}(f)$ (in [erg s^{-1} Hz^{-1}]), performed with an antenna telescope of effective area A_{eff}. We consider an unpolarised signal where the electric field vector changes randomly in all possible directions perpendicular to the direction of wave propagation. Since each of the two orthogonal polarisation directions contributes half the total flux, the power received from the source per unit frequency bandwidth and polarisation is $P_{f,source} = S_{f,obs}(f)A_{eff}/2$. Now consider modelling the antenna as a simple circuit consisting of a resistive load at temperature T. The thermal motions of the electrons inside the resistor induce a time variable voltage in this circuit which produces a power per unit frequency bandwidth that can be translated into a temperature by $P_{f,circuit} = k_B T$, where k_B is the Boltzmann constant. This is also known as Johnson–Nyquist noise (Johnson, 1928; Nyquist, 1928). Therefore an incident electromagnetic wave with a power per unit frequency $P_{f,source}$ will produce a power in the circuit given by $P_{f,circuit} = P_{f,source}$. It follows that:

$$S_{f,obs}(f) = \frac{2k_B T_{source}}{A_{eff}} = \frac{T_{source}}{G},$$ (1.97)

where we define the antenna gain as $G \equiv A_{\mathrm{eff}}/(2k_{\mathrm{B}})$ which is a measure of the antenna sensitivity. In other words the source of flux density $S_{f,\mathrm{obs}}(f)$ have an equivalent temperature T_{source} given by Equation (1.97) in the telescope system.

However real observing systems have many sources of noise, that can be modelled as a background system temperature T_{noise}. This noise temperature accounts for the noise contributions from the receiver, transmission lines, and other components of the radio telescope system which are represented as a system temperature T_{sys}. Furthermore, T_{noise} includes any spurious background signals coming from the sky with an equivalent temperature $T_{\mathrm{sky}}(l,b)$ which is a function of galactic sky coordinates (l,b). Figure 1.11 shows the Sky temperature distribution measured at 408 MHz as a function of equatorial coordinates (Haslam et al., 1982; Remazeilles et al., 2015). This temperature can be re-scaled to the central frequency of the survey assuming a frequency spectrum $T_{\mathrm{sky}} \propto f^{-2.6}$ (Lawson et al., 1987; Johnston et al., 1992). Therefore in general we can write $T_{\mathrm{noise}} = T_{\mathrm{sys}} + T_{\mathrm{sky}}(l,b)$. These noise sources are usually much larger than the faint signals we receive from astronomical sources. For this reason, to establish if a source is detectable we rely on the concept of signal-to-noise ratio which measures deviations from the random noise oscillation of the receiver.

The root mean square fluctuations σ_T of the noise in the telescope for a given generic temperature T decreases with the square root of the number of measured polarisations n_{pol}, the receiver frequency bandwidth Δf_{bw} and integration time t_{obs} (Dicke, 1946):

$$\sigma_T = \frac{T}{\sqrt{n_{\mathrm{pol}}\Delta f_{\mathrm{bw}}t_{\mathrm{obs}}}}. \tag{1.98}$$

This equation is known as the radiometer equation and is used to perform estimation of the system sensitivity.

Consider an observation of a pulsar for a total integration time t_{obs}. The pulsar has a top-hat pulse shape with width w_{obs}, periodicity P and an associated temperature T_{source}. Over the whole integration time, the pulsar is on over a time $t_{\mathrm{on}} = t_{\mathrm{obs}}w_{\mathrm{obs}}/P$ and it is off for a time $t_{\mathrm{off}} = t_{\mathrm{obs}}(P - w_{\mathrm{obs}})/P$. For the pulsar to be detectable, its signal has to exceed the noise fluctuations in the system during the integration time. The ratio between T_{source} and the system noise fluctuations $\sigma_{T_{\mathrm{noise}}}$ define the signal to noise ratios S/N. To have a significative detection the S/N has to surpass a threshold that is usually taken to be above 5 (see Vivekanand, Narayan, and Radhakrishnan, 1982; Lorimer and Kramer, 2012). The root mean square noise fluctuations during the integration time can be computed as the quadrature sum:

$$
\begin{aligned}
\sigma_{T_{\mathrm{noise}}} &= \sqrt{\sigma_{T_{\mathrm{on}}}^2 + \sigma_{T_{\mathrm{off}}}^2} \\
&= \sqrt{\left(\frac{T_{\mathrm{source}} + T_{\mathrm{noise}}}{\sqrt{n_{\mathrm{pol}}\Delta f_{\mathrm{bw}}t_{\mathrm{on}}}}\right)^2 + \left(\frac{T_{\mathrm{noise}}}{\sqrt{n_{\mathrm{pol}}\Delta f_{\mathrm{bw}}t_{\mathrm{off}}}}\right)^2} \\
&\simeq \frac{T_{\mathrm{noise}}}{\sqrt{n_{\mathrm{pol}}\Delta f_{\mathrm{bw}}t_{\mathrm{obs}}}}\frac{P}{\sqrt{w_{\mathrm{obs}}(P - w_{\mathrm{obs}})}},
\end{aligned}
\tag{1.99}
$$

where we used Equation (1.98) and made the approximation that $T_{\text{source}} \ll T_{\text{noise}}$. We can therefore define the signal-to-noise ratio as:

$$
\begin{aligned}
S/N &\equiv \frac{T_{\text{source}}}{\sigma_{T_{\text{noise}}}} = \frac{T_{\text{source}}\sqrt{n_{\text{pol}}\Delta f_{\text{bw}}t_{\text{obs}}}}{T_{\text{noise}}}\frac{\sqrt{w_{\text{obs}}(P - w_{\text{obs}})}}{P} \\
&= \frac{S_{\text{obs}}(f)\sqrt{n_{\text{pol}}\Delta f_{\text{bw}}t_{\text{obs}}}}{G[T_{\text{sys}} + T_{\text{sky}}(l,b)]}\frac{\sqrt{w_{\text{obs}}(P - w_{\text{obs}})}}{P},
\end{aligned}
\tag{1.100}
$$

where in the last step we used Equation (1.97). To take into account system imperfections, for example due to the digitisation of the signal, a degradation factor $\beta \sim 1.5$ is commonly introduced so that the previous equation becomes:

$$
S/N = \frac{S_{f,\text{obs}}(f)\sqrt{n_{\text{pol}}\Delta f_{\text{bw}}t_{\text{obs}}}}{\beta G[T_{\text{sys}} + T_{\text{sky}}(l,b)]}\frac{\sqrt{w_{\text{obs}}(P - w_{\text{obs}})}}{P}.
\tag{1.101}
$$

If we consider the period-averaged flux by substituting Equation (1.91) into (1.101) we find:

$$
S/N = \frac{S_{f,\text{mean}}(f)\sqrt{n_{\text{pol}}\Delta f_{\text{bw}}t_{\text{obs}}}}{\beta G[T_{\text{sys}} + T_{\text{sky}}(l,b)]}\sqrt{\frac{P - w_{\text{obs}}}{w_{\text{obs}}}}.
\tag{1.102}
$$

This equation shows that the narrower the pulse, the bigger the signal-to-noise ratio, hence the easier it is to detect the pulsar.

1.13 The neutron star zoo

Figure 1.12 shows the observed population of non-accreting neutron stars in the $P - \dot{P}$ plane (where P is the neutron star spin period and $\dot{P}$ is the spin period derivative with respect to time). Assuming a dipolar spin-down mechanism we also overlay the lines of constant polar magnetic field at the surface (*solid lines*) (see Equation (1.46)) and characteristic age (*dashed lines*) (see Equation (1.48)) as a reference. Around ~ 3000 neutron stars are known to date and according to their observational properties we classify them as different types. This section will provide an overview of the different classes of neutron stars and their distinguishing characteristics (see Kaspi, 2010; Harding, 2013, for a more exhaustive review).

Rotation-powered pulsars. The bulk of the population is composed of rotation-powered pulsars (RPPs) (*grey dots*). As they spin down due to magnetospheric torques, the loss of rotational energy is converted into particle acceleration which in turn produces broadband electromagnetic emission, from radio to X-rays and gamma-rays. The origins of the radio and gamma-ray emission of pulsars are however not well understood. As described in Section 1.9, the radio emission is thought to be generated close to the neutron star surface in a narrow beam centred on the magnetic axis within the magnetic pole region of the pulsar. The high-energy gamma-ray emission is instead thought to be generated in the outer regions of the neutron star magnetosphere close to the light cylinder where the magnetic field lines open up (see also Section 1.7). Since the discovery of the first neutron star through its radio pulsations (Hewish et al., 1968) around ~ 2000 radio pulsars have been detected. For some of them, X-ray, gamma-ray and optical pulsations were soon discovered.
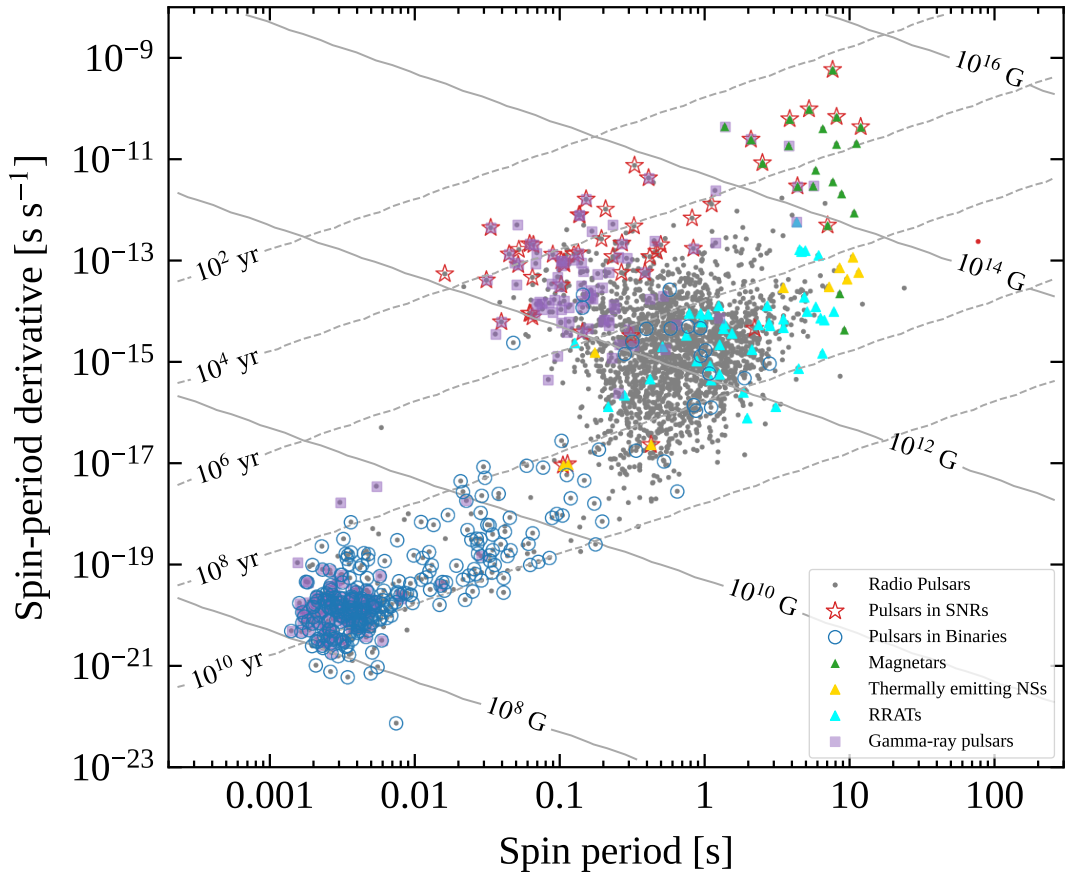
FIGURE 1.12: $P - \dot{P}$ diagram showing the spin period derivative *vs* spin period for the detected neutron stars (Data from the ATNF Pulsar Catalog, https://www.atnf.csiro.au/research/pulsar/psrcat/; Manchester et al., 2005).

At present, there are over 100 RPPs detected at X-ray energies and over 300 gamma-ray pulsars (Abdo et al., 2013; Smith et al., 2023). While most were found using the ephemerides of already known radio pulsars, many were discovered through their X-ray or gamma-ray pulsations and are radio-quiet.

Young neutron stars populate the top part of the diagram. They are often associated with Supernova Remnants (SNRs) (*red stars*) and are often also detected as high-energy gamma-ray pulsars (*purple squares*). As they spin down and their magnetic fields decay in time (see Section 1.6, pulsars move towards the bottom right region of the diagram. There are two main populations of RPPs. "Normal" pulsars with spin periods $P \gtrsim 0.1\,\text{s}$ and characteristic ages $\tau \lesssim 100\,\text{Myr}$, occupy the central part of the diagram whereas millisecond pulsars (MSPs), with spin periods down to a few milliseconds and generally higher characteristic ages $\tau \gtrsim 100\,\text{Myr}$ stretch towards the bottom left part of the diagram. MSPs are believed to be older neutron stars mostly in binary systems. Originally part of the X-ray binary population, they are recycled, i.e., spun up during a long phase of accretion from their binary companion (Bhattacharya and van den Heuvel, 1991; Tauris and Savonije, 1999; Tauris, Langer, and Kramer, 2012). There is another subpopulation of RPPs known as rotating radio transients (RRATs) (*light blue triangles*). They were discovered as powerful short-lasting radio bursts that seemed to recur randomly over timescales of minutes to hours (McLaughlin et al., 2006). By studying the arrival times of these bursts common periodicities have been discovered hinting to the spin periods of the emitting neutron stars. RRATs have been interpreted as normal pulsars showing an irregular, sporadic emission of radio pulses (detected as radio bursts), instead of a regular emission associated with the rotation period. For these characteristics they are difficult to detect but they should be common in the neutron star population.

Magnetars.   Neutron star emission can also be powered by magnetic energy. This is the case for the so-called magnetars, historically also known as soft gamma-ray repeaters (SGRs) and anomalous X-ray pulsars (AXPs) (Woods and Thompson, 2006), which populate the upper-right corner of the $P - \dot{P}$ diagram (*green triangles*). Magnetars have incredibly strong magnetic fields estimated to exceed $10^{14}\,\text{G}$. This immense magnetic energy reservoirs often result in sporadic bursts of X-rays and gamma-rays which are thought to be caused by the magnetic field decay and significant magnetic field readjustments in the star's crust, leading to starquakes and intense releases of energy (see Kaspi and Beloborodov, 2017; Esposito, Rea, and Israel, 2021, for reviews).

Thermally emitting neutron stars.   Some neutron stars emit thermal radiation, primarily in the form of X-rays, due to their high surface temperatures. These objects are collectively referred to as thermally emitting neutron stars (*yellow triangles*). They include two main subclasses: central compact objects (CCOs) and the X-ray dim isolated neutron stars (XDINSs). CCOs are found near the centres of supernova remnants and therefore should be young neutron stars with ages not exceeding a few tens of kyrs. However, they have low $\dot{P}$ values which translate into a low inferred magnetic fields in the dipole braking scenario. A possible explanation for this could be that their fields have been buried during an episode of fallback accretion of debris after the supernova explosion (Shabaltas and Lai, 2012; Viganò and Pons, 2012; Torres-Forné et al., 2016). XDINSs are also called the magnificent seven and are among the closest neutron stars that we know (few hundred parsecs from the Earth) (see van Kerkwijk and Kaplan, 2007; Turolla, 2009, for reviews). Only six of

them have measured timing properties. Their thermal emission comes directly from the neutron star surface and is mainly powered by the residual heat stored in their interior and by Ohmic dissipation of the magnetic field in the neutron star crust. Therefore the study of thermally emitting neutron stars provides valuable insights into the cooling mechanisms and interior composition of neutron stars (Viganò et al., 2013).

The strict distinction between these different neutron star classes has started to be questioned after some recent discoveries. For example, the dichotomy between rotational-powered pulsars and magnetars has been softened because magnetar-like X-ray activity was found in pulsars with large rotational energy loss rates, such as PSR J1846–0258 (Gavriil et al., 2008) and PSR J1119–6127 (Archibald et al., 2016). Moreover pulsed radio emission has now been detected from several magnetars in outburst like XTE J1810–197 (Camilo et al., 2006), AXP J1810–197 (Kramer et al., 2007) and J1818.0–1607 (Esposito et al., 2020).

Furthermore recent radio surveys, in particular thanks to new radio interferometers such as the LOw Frequency ARray (LOFAR; van Haarlem et al., 2013), MeerKAT (Jonas, 2009), The Australian SKA Pathfinder (ASKAP; Hotan et al., 2021), and the Murchison Widefield Array (MWA; Tingay et al., 2013; Wayth et al., 2018), have started to uncover the existence of a new population of mysterious radio sources with very long period that challenge our understanding of the pulsar population and its evolution.

1.13.1   Long-period pulsars

Until recently, the only neutron star known with a very long spin period was the X-ray emitting neutron star 1E 161348-5055 at the centre of the 2 kyr-old supernova remnant RCW103. This source manifests a measured modulation of ~ 6.67 hr, and was responsible of a large magnetar-like X-ray outburst (Rea et al., 2016; D'Aì et al., 2016) which demonstrated the source's isolated magnetar nature despite its long period and young age.

However, in the last couple of years a ~ 76 s radio pulsar (PSR J0901-4046, Caleb et al., 2022) and two peculiar coherent radio sources with a periodicity of ~ 1091 s (GLEAM-X J162759.5-523504.3; Hurley-Walker et al., 2022a) and ~ 1318 s (GPM J1839–10 Hurley-Walker et al., 2023) have been discovered. The nature of these two latter sources is still debated. Their emission properties such as the variable and irregular pulse profile, the high level of polarisation and the estimated luminosity exceeding their rotational power seem to point to a magnetar origin. Another proposed interpretation is that these long-period sources are magnetised white dwarfs with an active magnetosphere able to emit pulsar-like coherent radio emission.

We will explore possible scenarios to produce long period pulsars, by considering spin-down induced on a neutron star by the interaction with supernova fallback matter (see Chapter 3) and performing population synthesis studies of neutron stars and magnetic white dwarfs (see Chapter 4). These discoveries overall are showing that the boundaries between the different neutron star classes are blurring and motivate theoretical frameworks that consider the different neutron star types in an evolutionary scenario.

1.14 The birth rate and the need for an evolutionary scenario

As neutron stars are born from the core collapse of massive OB stars (see Section 1.2), their birth rate should not exceed the expected core-collapse supernova (CCSN) rate in the Milky Way. The Galactic CCSN rate can be estimated via different methods:

- Reed (2005) performed a census of the massive stars (with mass $M \gtrsim 10 \, M_\odot$) in the solar neighbourhood. From this sample, they extrapolated a Galactic CCSN rate between 1 and 2 supernovae per century.

- Li et al. (2011) studied the CCSNe in the local Universe and by assuming that the Milky Way has similar properties as other galaxies with the same morphological type they derived a CCSN rate of 2.30 ± 0.48 per century.

- Diehl et al. (2006) modelled the gamma-ray emission from radioactive ^{26}Al in the Milky Way by assuming that this emission traces the ongoing nucleosynthesis pollution by CCSNe in the Milky Way. They estimated in this way a CCSN rate of 1.9 ± 1.1 per century in our Galaxy.

By combining these different ways of estimating the Galactic CCSN rate, a recent study from Rozwadowska, Vissani, and Cappellaro (2021) found a best estimate of $\mathcal{R}_{\mathrm{CCSN}} = 1.79 \pm 0.55$ per century. The birth rate of Galactic neutron star should therefore be compatible with this estimate and not exceed this value.

In the past, several works have used the population of radio pulsars to make estimates of their birth rate. Historically, pulsar population analyses used the fact that the typical timescales on which neutron stars are active as radio pulsars are much shorter than the age of the Galaxy. In this way, the population can be considered in a steady state and can be viewed as "flowing" in the $P - \dot{P}$ diagram as a "current" obeying a continuity equation (Phinney and Blandford, 1981; Vivekanand and Narayan, 1981). At a given spin period P', this current, $J(P')$, equals the pulsar birth rate, $\mathcal{R}_{\mathrm{birth}}$, minus the pulsar death rate, $\mathcal{R}_{\mathrm{death}}$, in the period range from zero to the actual spin period P', i.e.:

$$J(P') = \mathcal{R}_{\mathrm{birth}}(0 < P < P') - \mathcal{R}_{\mathrm{death}}(0 < P < P'). \tag{1.103}$$

As the known radio pulsar population is flux-limited (depending on the discovery survey), the maximum value of $J(P)$ provides a lower limit on the pulsar birth rate. For example, Lorimer et al. (2006), using 1008 non-recycled pulsars detected by the Parkes Low-Latitude Survey estimated a birth rate of 1.38 ± 0.21 per century with this method. The advantage of this approach to compute birth rates is that it is agnostic to the details of how pulsars evolve in time.

A second approach to estimate pulsar birth rates is using population synthesis, discussed in detail below. Contrary to the previous approach, population synthesis relies on assuming physical models to describe pulsar properties and how the pulsar population evolves in time. For example, Faucher-Giguère and Kaspi (2006) use Monte Carlo simulations to model the birth properties of pulsars (velocity distributions, spin periods, magnetic fields) and the detectability in the Parkes and Parkes Swinburne Multibeam radio pulsar surveys. They evolve the initial population in time to obtain an observed synthetic pulsar sample to compare with observational data. In this way, they obtain a higher pulsar birth rate of 2.8 ± 0.5 per century compared to Lorimer et al. (2006).

We note that even without considering the existence of other neutron star classes, we are already at the limit of being compatible with the CCSN rate. Keane and

Kramer (2008) highlight that if one takes into account the birth rates estimated for other classes of neutron stars like RRATs, XDINSs and magnetars (see Section 1.13), the total neutron star birth rate would exceed the Galactic CCSN rate by a factor of $\sim 5 - 10$. Furthermore, the recent discovery of very long-period radio sources (see Chapters 4 and 5) will further increase this discrepancy if their origin is associated with neutron stars. A possible solution is to consider an evolutionary scenario that introduce relationships between the different neutron star classes. Indeed by taking into account the coupled evolution of the magnetic field and the temperature in neutron stars, Viganò et al. (2013) explain the emission from magnetars, XDINs, CCOs and the X-ray emitting rotational powered pulsars as manifestations of neutron stars of different ages and with different magnetic-field strengths and geometries.

1.15   Population synthesis: current state of the art

Population synthesis models aim to reproduce the observed population of neutron stars by simulating their formation, evolution, and observable properties. Typically the adopted approach relies on Monte Carlo simulations to model the different properties of neutron stars from some parametric distributions. These properties are then evolved in time through analytical and numerical models. Finally, models for the observational biases and filters are applied to assess which neutron stars can be detected to build a synthetic observed population. The simulation outcomes are then compared to the empirical observed data to assess the validity of theoretical models, constrain model parameters and gain insights into the underlying physics.

Due to the abundance of data in the radio band, many works on population synthesis have mainly focused on reproducing the radio pulsar population, in particular trying to constrain the birth spin-period and magnetic-field distributions. Some notable research papers in this area include Narayan and Ostriker (1990), Faucher-Giguère and Kaspi (2006), Gonthier et al. (2007), Kiel et al. (2008), Kiel and Hurley (2009), Popov et al. (2010), Osłowski et al. (2011), Levin et al. (2013), Gullón et al. (2014), Bates et al. (2014), and Cieślar, Bulik, and Osłowski (2020). Overall these works adopt various prescriptions for the physical models and techniques to perform parameter optimisation. In Table 1.1 I summarise the main prescriptions used by some reference works and compare them with the present work (see Chapters 5 and 6). For example, Narayan and Ostriker (1990) and Faucher-Giguère and Kaspi (2006) used the χ^2 metric and Kolmogorov-Smirnov (KS) tests to compare the simulated distributions with the observed ones and guide a "by hand" parameter optimisation. These methods manifest some limitations as the computation of χ^2-statistics relies on having a good amount of data while the KS-test is not particularly sensitive when applied to multi-dimensional distributions of data. In recent years, improved modelling of neutron star magnetic field evolution and decay, refined cooling models, increased amounts of empirical data and the development of Bayesian techniques like Markov Chain Monte Carlo (MCMC) have all contributed to refining population synthesis codes. Current research efforts aim to combine the information coming from multi-wavelength observations, to better constrain the properties of the neutron star population as a whole and refine theoretical models. A first attempt to include a more realistic magnetic field evolution based on magneto-thermal simulations of neutron stars and to consider both the radio and the X-ray pulsar population in a population synthesis study was performed by Popov et al. (2010), Gullón et al. (2014), and Gullón et al. (2015). In particular Gullón et al. (2014) and Gullón et al. (2015) employed annealing methods for optimisation but found degeneracies

between the simulation parameters were found and it was difficult to put meaningful constraints on the birth rate, especially for magnetars. Johnston et al. (2020) performed a population synthesis study combining both radio and gamma-ray pulsars and showed that highly energetic gamma-ray pulsars and young radio pulsars might come from the same underlying population. Cieślar, Bulik, and Osłowski (2020) reproduce the observed radio pulsar population using MCMC, although using a simulation framework with simplified physical models.

This thesis work represents the first efforts to developing a population synthesis approach for isolated neutron stars in conjunction with modern deep-learning techniques. Indeed deep learning can be a powerful tool for extracting relevant features from multi-dimensional data coming from different surveys and performing parameter inference (Chapter 2). In this work, we will focus on the radio pulsar population to benchmark our population synthesis approach and compare the results with previous ones in the literature. In particular, we will focus on predicting the birth properties, i.e., spatial and kick-velocity distributions (Chapter 5) and initial spin-period and magnetic-field distributions and magnetic-field evolution (Chapter 6). In Chapter 3 and 4 we will instead investigate the nature of recently discovered long-period radio sources assuming a neutron-star or white-dwarf origin scenarios.

We first turn our attention to machine learning, introducing the key concepts crucial for the following chapters.

TABLE 1.1: Comparison between the main prescriptions used in this work and some reference population synthesis studies in the literature.

	Faucher-Giguère and Kaspi (2006)	Bates et al. (2014)	Gullón et al. (2014) Gullón et al. (2015)	Cieślar, Bulik, and Osłowski (2020)	This work
Dynamics	kick velocity: exponential spiral arms, exponential disk	kick velocity: exponential, normal spiral arms, exponential disk	kick velocity: exponential, normal spiral arms, exponential disk	kick velocity: Maxwell spiral arms, exponential disk	kick velocity: Maxwell e-density distribution, exponential disk Galactic rotation
Birth spin-period distribution	normal	normal, log-normal	normal	normal	log-normal
Birth magnetic-field distribution	log-normal	log-normal	log-normal	log-normal	log-normal
Magnetic-field evolution	no B-decay	no B-decay	B-decay (magneto-thermal models)	exponential B-decay	B-decay (magneto-thermal models)
Rotational evolution	vacuum dipolar spin-down no alignment	vacuum dipolar spin-down exponential alignment	plasma-filled dipolar spin-down coupled spin-period, alignment evolution	vacuum dipolar spin-down no alignment	plasma-filled dipolar spin-down coupled spin-period, alignment evolution
Radio-beam intercept	empirical beam fraction[a]	empirical beam fraction[a] geometry dependent[b]	empirical beam fraction[a]	empirical beam fraction[a]	geometry dependent[b]
Radio luminosity	pseudo-luminosity $\propto \langle \dot{E}_{\rm rot}\rangle^\epsilon$	pseudo-luminosity $\propto P^\alpha \dot{P}^\beta$	pseudo-luminosity $\propto \langle \dot{E}_{\rm rot}\rangle^\epsilon$	pseudo-luminosity $\propto \langle \dot{E}_{\rm rot}\rangle^\epsilon$	intrinsic luminosity $\propto \langle \dot{E}_{\rm rot}\rangle^\epsilon$
Surveys	PMPS, SMPS	PMPS, SMPS	PMPS, SMPS + X-ray NS	PMPS	PMPS, SMPS, HTRU
Optimisation method	KS-test, by eye	KS-test	annealing method, KS-test	MCMC, Gaussian likelihood	Deep learning

[a] The pulsar intercepting our line of sight are determined using an empirical relation between the beam fraction and the spin period computed from polarisation data (see Tauris and Manchester, 1998).
[b] The pulsar intercepting our line of sight are determined considering the geometry of the radio beam aperture and the inclination angle of the magnetic dipole moment (see Chapter 6).

Chapter 2

Machine learning

2.1 Introduction

In the past decade, the accumulation of extensive and heavy datasets has been almost ubiquitous in astronomy and astrophysics. In order to take full advantage of these data and perform data-driven science that complements ongoing theoretical modelling efforts, new techniques and analysis pipelines that can handle these large amounts of data (and do so in an automated way) are required. Machine learning (ML) has played an important role in developing such new algorithms (Ball and Brunner, 2010; Allen et al., 2019; Baron, 2019; Fluke and Jacobs, 2020). As a subfield of artificial intelligence, machine learning involves the development of algorithms that allow computers to learn from and make predictions or decisions based on data without being explicitly programmed.

Depending on the task at hand, there are two different approaches to train machine-learning algorithms. The first one is called *supervised learning* and involves training a machine-learning model using labelled training data. In this approach, the dataset consists of input variables (features) and their corresponding output variables (labels). The goal is to learn a mapping function that predicts the labels for new, unseen data. In general to acquire good performance, a supervised learning algorithm requires a large amount of labelled data. Depending on the task, these algorithms can be categorised into two main types:

- **Classification**: This type of algorithm aims to predict discrete class labels for the given input data. An example is the prediction of the morphological type of a galaxy based on its image (e.g. Dieleman, Willett, and Dambre, 2015).

- **Regression**: The goal of these algorithms is to predict continuous numerical values based on the input features. An example is the prediction of the star formation rate of a galaxy based on its redshift, luminosity and colour (e.g. Bonjean et al., 2019).

Among the most common supervised learning algorithms, are decision trees (Quinlan, 1986), support vector machines (Burges, 1998), and artificial neural networks (Bishop, 1995).

The second approach is *unsupervised learning*. In contrast to supervised learning, unsupervised learning involves training a model on unlabelled data. Here, the dataset only consists of input data without any corresponding output labels. The objective of unsupervised learning is to find patterns, structures, or relationships within the data. In principle unsupervised learning algorithms can potentially work with smaller amount of data. They can be categorised into two main types:

- **Clustering**: These algorithms aim to group similar data points based on their inherent similarities or distances in a given representation space. The goal

is to discover hidden patterns or groupings within the dataset. For example Sánchez Almeida and Allende Prieto (2013) used clustering to classify stellar spectra.

- **Dimensionality reduction**: These algorithms seek to reduce the number of input variables by transforming the data into a lower-dimensional representation. This helps in visualising and compressing complex datasets by focusing on their most relevant properties. Usually these techniques are used as a pre-processing step. For example they have been employed to extract physical parameters of stellar atmospheres from their spectra (e.g. Re Fiorentin et al., 2007).

Examples of unsupervised learning algorithms include k-means clustering (MacQueen, 1967), principal component analysis (PCA, Hotelling, 1933), and isomap embedding (Tenenbaum, Silva, and Langford, 2000).

In the following, for the purpose of this thesis, I will focus on supervised machine-learning techniques using artificial neural networks to perform regression tasks.

2.2 Deep learning

Deep learning is a subset of machine learning that uses artificial neural networks formed of multiple layers of interconnected nodes (see Goodfellow, Bengio, and Courville, 2016, for a review). In analogy with biological brains, each node can be thought of as a neuron while the connections between nodes are the synapses. The importance of a connection is defined by a weight, w, represented by a real number. Each node receives weighted signals from other connected neurons, processes them and transmits the processed signal to other connected nodes. In this way, artificial neural networks are able to process a given input signal, x, by decomposing it into more complex features and produce a final output, y. Therefore, they can be regarded as a function $y = F_w(x)$ where the set of weights, w, is adjusted to produce a desired output, y, given an input, x. The employment of deep neural networks enables the extraction of intricate features and representations from the input data, leading to refined predictions for classification or regression tasks. In the following, I will describe several basic network architectures that are of concern for this work but that can also be used for a great variety of problems. I will specifically focus on the multi-layer perceptron and the convolutional neural network. In preparation for the following chapters, I will also describe how the training process works (Section 2.3).

2.2.1 Multi-layer perceptron

The multi-layer perceptron (MLP) was originally conceived as a binary classifier (Rosenblatt, 1958; Minsky and Papert, 1969) but the name today is also used to denote modern so-called fully-connected, deep, feedforward neural networks. In this type of architecture, the neurons are organised in layers and the information flows unidirectionally from the first to the last layer (see top sketch in Figure 2.1). The first layer is denoted as the input layer, the last one as the output layer and the ones in the middle as hidden layers. The nodes in a given layer are not connected and work in parallel, but the layers are fully connected to each other meaning that each neuron receives inputs from all nodes in the previous layer and passes its output to all nodes in the subsequent layer. The i-th neuron in a given layer performs a linear combination of the array of signals, x, coming from the previous layer according
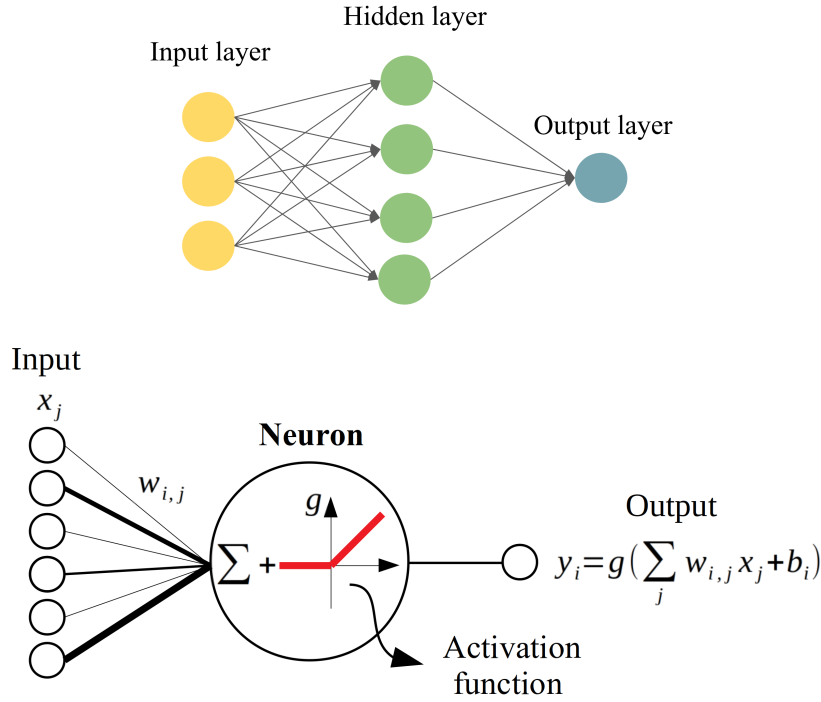
FIGURE 2.1: Sketch of a simple fully-connected feedforward neural network with one hidden layer (top panel) and of a neuron in a multi-layer perceptron (bottom panel). To produce an output y_i, each neuron first perform a linear combination of the input values x_j and then applies a non linear activation function g.

to the weights, $\boldsymbol{w}_i$, associated with its connections and a bias term, b_i (see bottom sketch in Figure 2.1):

$$z_i = \boldsymbol{w}_i^{\mathrm{T}} \boldsymbol{x} + b_i = \sum_{j=1}^{n} w_{i,j} x_j + b_i, \tag{2.1}$$

where n is the number of neurons in the previous layer and the superscript "T" denotes the transposed vector. The bias term is responsible of introducing an offset in the output value. By adding an extra term to the input $\boldsymbol{x}$ and setting it to 1, we can write the linear combination above in a more compact form:

$$z_i = \boldsymbol{w}_i^{\mathrm{T}} \boldsymbol{x} = \sum_{j=0}^{n} w_{i,j} x_j, \tag{2.2}$$

where we adopt the convention that $x_0 = 1$, so that $b_i = w_{i,0} x_0$ becomes the bias term. In the following, I will also use the term weight to refer to the bias terms for simplicity. Before passing the processed signal to all the neurons in the following layer a non-linear activation function, g, is usually applied to produce an output $y_i = g(z_i)$ (see Figure 2.1). Non-linear activation functions are used in neural networks to introduce non-linearity into the model (see Section 2.2.3). This is crucial to enable the neural network to learn complex, non-linear patterns and relationships between input and output variables. Therefore, in an MLP with a total of L layers, each layer represents a function $F^{(l)}$ with $l = 1, \dots L$ that performs a linear combination of the input and applies a non-linear function g. Therefore by passing through the network

the input signal x undergoes a composition of functions until the final output, y is computed:

$$y = F^{(L)}(F^{(L-1)}...(F^{(2)}(F^{(1)}(x)))) \tag{2.3}$$

The output layer will give back the prediction for either the regression or classification task, depending on the specific problem.

2.2.2   Convolutional neural network

Convolutional neural networks (CNNs) (LeCun et al., 1989) are deep-learning algorithms primarily used for computer vision tasks like image classification, object detection, and segmentation (Krizhevsky, Sutskever, and Hinton, 2012). They take inspiration from the visual cortex of living organisms and are designed to work with structured, grid-like data that are usually referred to as feature maps.

CNNs are composed of multiple layers, including so-called *convolutional layers*, pooling layers, and fully-connected layers. The key component is the convolutional layer, which performs the convolution operation on the input feature map using a set of filters or kernels composed of learnable weights. For example, Let us consider a two-dimensional feature map as input x and a two-dimensional kernel w. We can define the convolution operation $w * x$ over a discretised grid as (see Chapter 9 in Goodfellow, Bengio, and Courville, 2016):

$$z_{i,j} = (w * x)_{i,j} = \sum_{m=-\infty}^{+\infty} \sum_{n=-\infty}^{+\infty} w_{m,n} x_{i-m, j-n}. \tag{2.4}$$

The result of this operation will be non-zero only where the feature map and the filter overlap. However, nearly all machine learning and deep learning libraries implement the cross-correlation operation $w \star x$ instead and improperly call it convolution. This is defined as:

$$z_{i,j} = (w \star x)_{i,j} = \sum_{m=-\infty}^{+\infty} \sum_{n=-\infty}^{+\infty} w_{m,n} x_{i+m, j+n}. \tag{2.5}$$

Note that the difference between Equation (2.4) and (2.5) consists only in the "$\pm$" signs in the indexes of x which flips the way we access the entries of the input feature map. For simplicity, we will continue calling both operations convolution. In practice, the convolution operation involves sliding the filter over the input feature map, computing element-wise multiplications, and summing them up to produce a new convolution feature map z (see Figure 2.2). These feature maps capture local patterns or features present in the input map such as edges or corners in an image. Generally, the convolutional filters have a much smaller size if compared with the input feature map. By using small filters that slide over the input map each output value in the convolutional feature map depends only on a small subset of input values. This reduces the number of computations required compared to fully-connected layers where each input value affects all output values. Indeed fully-connected deep neural networks are not suitable to process multi-dimensional inputs, since the number of required neurons and connections will scale up exponentially with the input size. Furthermore, the same filter (set of weights) is used across different spatial locations of the input feature map. This sharing of parameters significantly reduces the number of parameters in the network and enables the model to learn spatially invariant features. By sharing parameters, the network can detect the same feature regardless
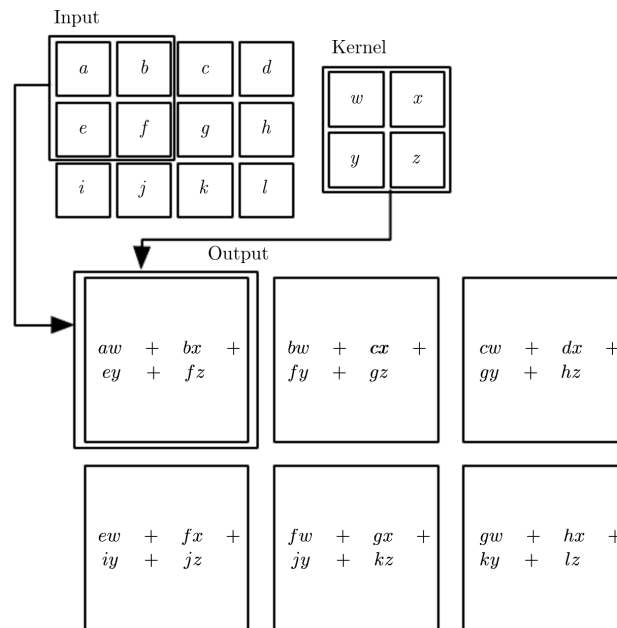
FIGURE 2.2: An example showing how the convolution (cross-correlation in this case) operation works for a 2-D input. We draw boxes with arrows to indicate how the upper-left element of the output feature map is formed by applying the kernel to the corresponding upper-left region of the input feature map (Figure taken from Goodfellow, Bengio, and Courville, 2016).

of its location and orientation in the input. Since a single filter can recognise only a specific type of feature, usually many filters are stacked together to form a convolution layer. This allows one to perform many convolution operations on the input map and to extract different features in parallel.

To introduce non-linearity, an activation function is usually applied to the feature maps generated after the application of a convolutional layer.

After the activation function has been applied to the feature maps, a common practice is to use *pooling layers*. The pooling operation performed by this kind of layers is specified, rather than learned. Two common functions used in the pooling operation are max-pooling and average-pooling. Max-pooling and average-pooling layers consist of filters that slide on the feature map and select the maximum or compute the average value respectively inside the sub-region of the map that they overlap with (see Figure 2.3). Both approaches are used to downsample the feature maps by reducing their spatial dimensions, thereby decreasing the computational complexity while preserving the most relevant information and the dominant features. In this way, these layers help reduce the influence of noisy or irrelevant information from the input.

Usually, deep CNNs are formed by stacking together a series of convolutional layers composed of an increasing number of filters followed by pooling layers. This allows layers close to the input to learn low-level features (e.g. lines, edges, corners) and layers deeper in the model to learn high-order or more abstract features, like shapes or objects.

Finally, fully connected layers are added at the end of the CNN to perform the tasks of classification or regression. They take the high-level features extracted by the previous layers and produce the final output.
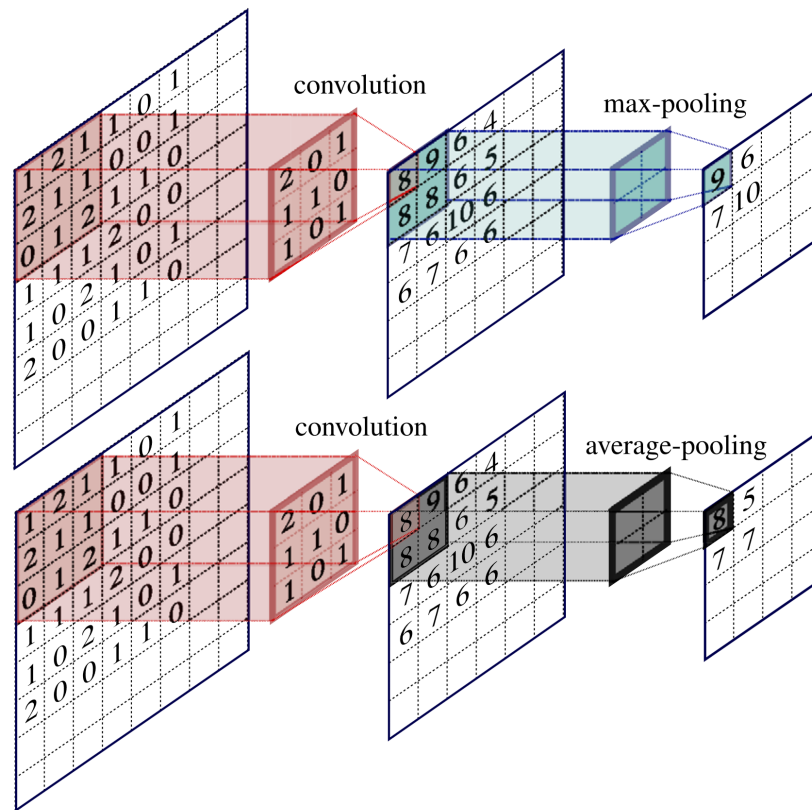
FIGURE 2.3: Example showing how the convolution and pooling layers operate on the input data. First, the 8×8 input layer is reduced to a 6×6 output layer after going through a convolution layer consisting of a single filter with size 3×3 that moves along the input with a stride of 1. This output layer becomes the input layer with respect to either the max-pooling or average-pooling layer, each consisting of a single filter with size 2×2 that moves along the input with stride 2. The max-pooling layer picks the maximum number from the layer input within a selected window, while the average-pooling layer computes average values over the selected window. The final output layer is therefore a 4×4 array (Figure taken from Teo et al., 2021).
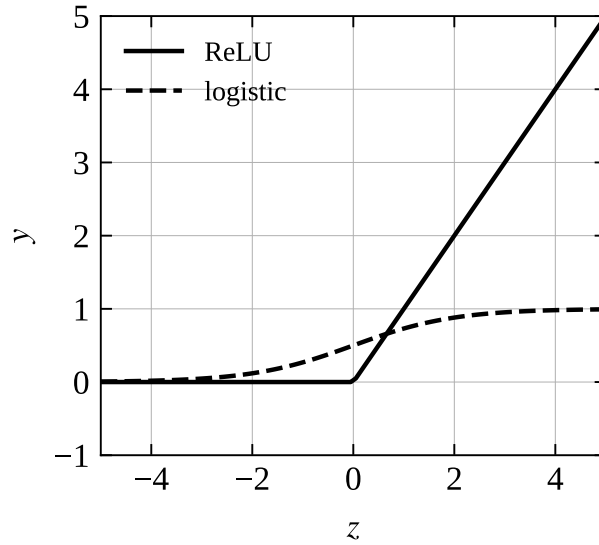
FIGURE 2.4: Comparison between the ReLU and the logistic (sigmoid) activation functions.

2.2.3 The activation function

Generally, the choice of the activation function depends on the problem and the type of data at hand. Besides introducing non-linearity the type of activation function is also important to implement a stable and effective learning process. An example of activation function g is the logistic (sigmoid) function:

$$\text{logistic}(z) = \frac{1}{1 + e^{-z}} \tag{2.6}$$

This activation function is especially used in the output layer of neural networks employed as classifiers as it compresses the output to a value between 0 and 1 which can be interpreted as the probability to belong to a given class. However as argued by Glorot and Bengio (2010) and Glorot, Bordes, and Bengio (2011) using the sigmoid function in all the layers of deep neural networks can slow down the learning process as it is very sensitive to the input value z especially when it is near 0. Moreover it saturates across most of its domain leading more easily to vanishing gradient issues (see Section 2.3 and Figure 2.4). Another common choice for the activation function is the rectified linear unit (ReLU) (Glorot and Bengio, 2010) which is defined as:

$$\text{ReLU}(z) = \max(0, z). \tag{2.7}$$

One major benefit of using ReLU is its simplicity (see Figure 2.4). As we will see in the next section Section 2.3, the gradient calculation for ReLU is usually easier and more efficient compared to more complex activation functions (it is 0 for $z < 0$ and always 1 for $z > 0$), resulting in faster convergence during gradient-based optimisation. Indeed for this activation function, there is a reduced chance of its gradient to vanish contrary to other activation functions such as the sigmoid. Furthermore by using the ReLU, some of the neurons are likely to be inactive (outputting zero) for negative input values. This sparse activation makes the network more efficient by reducing the number of active neurons and, in turn, the complexity and computational requirements (Glorot, Bordes, and Bengio, 2011).

2.3   The training process

In supervised learning, a neural network is trained on a dataset containing input features x and the corresponding target labels $\hat{y}$. By feeding the training data into the network, the training proceeds by adjusting the model parameters w, i.e., the weights and biases, to minimise a loss function. The loss is a measure of the difference between the predicted outputs y and the expected target labels $\hat{y}$. Its minimisation is achieved by computing the gradients of the loss function with respect to the weights and by performing gradient descent in a procedure called backpropagation (see Section 2.3.4). Here I will describe the steps involved in the training process.

2.3.1   Training, validation and test dataset

The labeled dataset that is available for training is usually divided into three separated subsets called training, validation, and test dataset, respectively. In machine learning, these are used for different purposes.

The training dataset is used to train a machine-learning model. It comprises a large portion (usually around 70% to 80%) of the original dataset. During the training phase, the model learns from this dataset by adjusting its parameters (i.e., the weights) to minimise the error between its predicted outputs and the actual target labels.

The validation dataset is used to evaluate the performance of the model during the training process. It usually constitutes a smaller portion (around 10% to 20%) of the original dataset. The validation dataset is not used to optimise the model parameters, but instead allows us to validate how well the model can generalise to unseen data. In general, it is used to fine-tune hyperparameters (like the learning rate, the size of input batches, etc., see below for an explanation), select the best model, and prevent overfitting, which refers to the model performing well on training data but poorly on unseen data.

Similar to the validation dataset, the test dataset is also a small fraction of the total dataset, usually around 10%. However, contrary to the validation dataset, the test dataset is never seen by the model during the training process and it is not used during model development or hyperparameter tuning. Instead it is employed to assess the final performance and generalisation ability of the trained model on real-world data in an unbiased way. To ensure this, the test dataset is completely separate from the training and validation datasets and it should ideally have a similar distribution of data.

Before training, it is sometimes useful to apply some transformations to these datasets. Some of these transformations include normalisation and standardisation of the input features x and in regression tasks also of the corresponding labels $\hat{y}$. In general, the label $\hat{y}$ can be a vector containing the different parameters to predict. Normalisation is usually applied on every input feature separately. In particular, it performs the following operation:

$$x' = \frac{x - \min{(x)}}{\max{(x)} - \min{(x)}}.$$
(2.8)

where $\min(x)$ and $\max(x)$ are the minimum and maximum values of the input feature map x respectively. This guarantees that the input values always fall in the range $[0, 1]$. The normalisation of the labels instead is performed on a dataset basis

and for each specific label $\hat{y}$ in $\hat{\boldsymbol{y}}$ it consists of the following operation:

$$\hat{y}' = \frac{\hat{y} - \min{(\hat{y})}}{\max{(\hat{y})} - \min{(\hat{y})}}, \tag{2.9}$$

where min and max are computed for each label parameter over the entire training dataset. In this way, all the labels are re-scaled to the range [0, 1] as well. Standardisation instead consists of evaluating the average value $\langle \boldsymbol{x} \rangle$ and the standard deviation $\sigma(\boldsymbol{x})$ for each input feature $\boldsymbol{x}$ and then computing the z-scores by applying the following transformation:

$$\boldsymbol{x}' = \frac{\boldsymbol{x} - \langle \boldsymbol{x} \rangle}{\sigma(\boldsymbol{x})}. \tag{2.10}$$

For the labels we standardise on a dataset basis and for each specific label $\hat{y}$ in $\hat{\boldsymbol{y}}$ we compute:

$$y'_{\text{t}} = \frac{\hat{y} - \langle \hat{y} \rangle}{\sigma(\hat{y})}, \tag{2.11}$$

where the mean and standard deviation are computed over the entire training dataset. Standardisation guarantees that the values of the input features and the labels are rescaled so that they have zero mean and unit variance.

Normalisation and standardisation aim to make the input features and targets have approximately the same range of values to improve the efficiency and effectiveness of the models' learning process. This is important because it prevents some variables from dominating others, which can lead to biased and ineffective learning. By normalising or standardising the input data, the model can therefore focus more accurately on the variations and patterns in the features rather than their absolute values. These preprocessing techniques facilitate effective learning by reducing the impact of outliers, and by providing a consistent and balanced representation of the input data.

2.3.2 Weights initialisation

An important aspect of achieving good performance while training is the initialisation of the weights of the network. When defining the network architecture, the simplest approach is to initialise all the weights randomly in a given range of values. In this way at the beginning of the training process also the output of the network is random. However, if the weight values are too small, after a few layers the output values and the gradients will completely vanish. On the other hand, if weights are too big the output values and the gradients could explode. As a consequence, training will proceed very slowly or will stagnate.

To visualise this aspect let us consider an example with a fully-connected neural network. We focus on two neighboring layers in the network, $l - 1$ and l, with layer $l - 1$ having n nodes. Suppose that the weights, w, associated to the connections between these two layers, are initialised with values randomly sampled from a normal distribution with mean 0 and a generic variance $\text{Var}(w)$ and the biases are initialised to be zero.

As shown in Equation (2.1), the outputs of the nodes in layer l are a linear combination of the input features x with weights w, i.e., output z is given by:

$$z^{(l)} = \sum_{i=1}^{n} w_i x_i. \tag{2.12}$$

The input features in turn are the result of the activation function applied to the output of the previous layer that is $x_i = y_i^{(l-1)} = g\left(z_i^{(l-1)}\right)$. We can calculate the expectation value of $z^{(l)}$:

$$\begin{aligned}
\mathrm{E}(z^{(l)}) &= \mathrm{E}\left(\sum_{i=1}^{n} w_i x_i\right) \\
&= n\mathrm{E}(wx) \\
&= n\mathrm{E}(w)\mathrm{E}(x) = 0,
\end{aligned} \tag{2.13}$$

where we treat the weights and the input features as two statistically independent random variables. Hence we dropped the index i for clarity of notation and we used the fact that the average value of the weights $\mathrm{E}(w) = 0$. Moreover for the variance of z we find:

$$\begin{aligned}
\mathrm{Var}(z^{(l)}) &= \mathrm{Var}\left(\sum_{i=1}^{n} w_i x_i\right) \\
&= n\mathrm{Var}(wx) \\
&= n[\mathrm{Var}(w)\mathrm{Var}(x) + \mathrm{E}(w)^2\mathrm{Var}(x) + \mathrm{Var}(w)\mathrm{E}(x)^2] \\
&= n[\mathrm{Var}(w)\mathrm{Var}(x) + \mathrm{Var}(w)\mathrm{E}(x)^2] \\
&= n\mathrm{Var}(w)[\mathrm{Var}(x) + \mathrm{E}(x)^2] \\
&= n\mathrm{Var}(w)\mathrm{E}(x^2),
\end{aligned} \tag{2.14}$$

where in the third line we used the rule for the variance of a product of random variables (Springer, 1979). In the fourth line we took advantage of the fact that we define the weights to have zero mean and, in the last line we use the definition of the variance $\mathrm{Var}(x) = \mathrm{E}(x^2) - \mathrm{E}(x)^2$. From this expression, we note that if the variance $\mathrm{Var}(w)$ and the expectation value $\mathrm{E}(x^2)$ are too big or too small, the variance $\mathrm{Var}\left(z^{(l)}\right)$ will explode or vanish after a few layers. Furthermore, it is worth noticing that as x is the result of applying the activation function to $z^{(l-1)}$, the expectation value and variance of x will, in general, differ from the ones of $z^{(l-1)}$ and they will depend on the type of activation function employed.

In the following we compute $\mathrm{E}(x^2)$ for the ReLU activation function, that is considering $x = \max\left(0, z^{(l-1)}\right)$ assuming that the probability density functions for the

variable x and z are $\mathcal{P}(x)$ and $\mathcal{P}(z)$ respectively:

$$
\begin{aligned}
\mathrm{E}(x^2) &= \int_{-\infty}^{+\infty} x^2 \mathcal{P}(x) dx \\
&= \int_{-\infty}^{+\infty} \max(0, z^{(l-1)})^2 \mathcal{P}(z) dz \\
&= \int_{0}^{+\infty} \left(z^{(l-1)}\right)^2 \mathcal{P}(z) dz \\
&= \frac{1}{2} \int_{-\infty}^{+\infty} \left(z^{(l-1)}\right)^2 \mathcal{P}(z) dz \\
&= \frac{1}{2} \mathrm{Var}\left(z^{(l-1)}\right),
\end{aligned}
\tag{2.15}
$$

where in the last step we recall that from Equation (2.13), $\mathrm{E}(z) = 0$ in any layer. Therefore, substituting in Equation (2.14) we find that:

$$
\mathrm{Var}\left(z^{(l)}\right) = \frac{1}{2} n \mathrm{Var}(w) \mathrm{Var}\left(z^{(l-1)}\right).
\tag{2.16}
$$

A good weight initialisation should ensure that $\mathrm{Var}\left(z^{(l)}\right) = \mathrm{Var}\left(z^{(l-1)}\right)$ which implies that the weights have a variance of:

$$
\mathrm{Var}(w) = \frac{2}{n^{(l-1)}}.
\tag{2.17}
$$

where we recall that $n^{(l-1)}$ is the number of nodes for the $l-1$ layer (i.e., the size of the input features for the l layer). Notice that in the previous steps we did not choose a specific layer l. Thus, this expression holds for every layer of the network. This is the initialisation procedure commonly adopted for models using the rectified activation function ReLU and is referred to as Kaiming initialisation (He et al., 2015). With this method, the initial weights between layers are sampled from a normal distribution with mean 0 and standard deviation $\sigma(w) = \sqrt{2/n}$, where n is the size of the input feature to each layer, while the biases are initially set to zero. This helps to reduce the problem of vanishing or exploding gradients and to make the training more stable and effective. In the following, we will adopt this initialisation procedure.

2.3.3 Forward pass and loss calculation

During the forward pass, input data from the training dataset is passed through the network, and intermediate activations of neurons are computed using the current weights. This process starts from the input layer and progresses through the hidden layers until it reaches the output layer. The output of the network is then compared to the desired output. The distance between the network's output y and the expected target label $\hat{y}$ is quantified using a loss function $\mathcal{L}(y, \hat{y})$, which measures the network's performance on a particular task. For regression tasks where one wants to predict some continuous real values, an example of a loss function is the mean squared error defined as:

$$
MSE(\boldsymbol{y}, \hat{\boldsymbol{y}}) = \frac{1}{n^{(o)}} \sum_{i=1}^{n^{(o)}} (y_i - \hat{y}_i)^2,
\tag{2.18}
$$
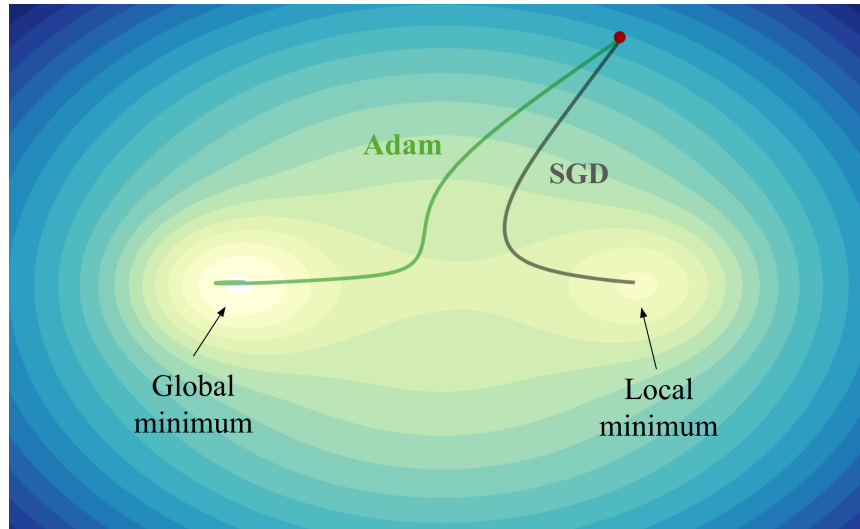
FIGURE 2.5: Visualisation showing the comparison in behaviour of the optimisation algorithms Adam (*green curve*) and stochastic gradient descent (SGD) (*dark grey curve*). The loss landscape has two minima, the global minimum is on the left while a local minimum is found on the right. Starting from the same point (*red dot*), Adam is able to find the global minimum while SGD converge to the local one (Figure adapted from https://emiliendupont.github.io/2018/01/24/optimization-visualization/).

where $n^{(o)}$ is the size of the output y.

Usually, the training dataset is passed through the network in batches composed of more than one sample. By passing a single training sample at a time, the gradient estimation can be noisy and the weight updates can be more erratic. On the other hand, using a larger batch size usually allows more stable training and gradient estimation since the loss computation and its gradients are averaged over the batch of data. This leads to a better estimate of the overall direction in which the weights need to be updated (see Section 2.3.4). The typical batch sizes are usually powers of 2, such as 32, 64, or 128. However, the appropriate batch size can vary depending on the dataset size, available computational resources, and the specific problem being addressed, and might need to be determined through experimentation and performance evaluation. Using a large batch size has indeed the downside of using more memory resources.

Once a batch of samples from the training dataset has been passed through the network and the loss has been computed the next step is the backpropagation.

2.3.4 Backward pass (backpropagation) and weight update

In this step, the loss is propagated backward through the network to compute the gradients with respect to the weights. Indeed the value of the loss function $\mathcal{L}(y, \hat{y})$ depends on the output y which in turn depends on the values of the weights in the network. The chain rule of calculus is applied to efficiently calculate the gradients of the loss with respect to each weight in the network. For example let us assume for simplicity that $y = g(z) = g(\sum_{i=0}^{n} w_i x_i)$ where g is again the activation function. We can now focus on a particular weight w_i. The gradient of the loss function with

respect to this weight is computed as:

$$\frac{\partial \mathcal{L}(y, \hat{y})}{\partial w_i} = \frac{\partial \mathcal{L}(y, \hat{y})}{\partial y} \frac{\partial y(z)}{\partial z} \frac{\partial z(\boldsymbol{w}, \boldsymbol{x})}{\partial w_i}. \tag{2.19}$$

We note here that the gradient of the loss function vanishes if either $\frac{\partial y}{\partial z}$ or $\frac{\partial z}{\partial w_i}$ or both are zero as could happen for an unsuitable choice of activation functions or a bad weight initialisation (see Section 2.2.3 and Section 2.3.2).

The gradients calculated for the backpropagation are then used to update the weights of the network. This is usually performed through gradient descent. Let us again focus on a specific weight w_i for simplicity.

- If $\frac{\partial \mathcal{L}}{\partial w_i} > 0$, an increase in the value of w_i increases the loss $\mathcal{L}$;

- If $\frac{\partial \mathcal{L}}{\partial w_i} < 0$, an increase in the value of w_i decreases $\mathcal{L}$.

Therefore to always decrease the value of the loss $\mathcal{L}$, the update δw_i in the weight value has to be proportional to the opposite of the gradient computed with respect to that weight. This means that, in general, at the training iteration t, the weight values can be updated in the following way:

$$\boldsymbol{w}_t = \boldsymbol{w}_{t-1} + \delta \boldsymbol{w}_t = \boldsymbol{w}_{t-1} - \eta \boldsymbol{\nabla}_w \mathcal{L}(\boldsymbol{w}_{t-1}), \tag{2.20}$$

where η is called the learning rate defined as $\eta > 0$. The learning rate establishes the size of the displacement between the updated values and the old values of the weights at training step $t - 1$.

In the simplest optimisation algorithm called stochastic gradient descent (SGD), the learning rate is kept constant. More sophisticated optimisation algorithms introduce a more efficient way to update the weights. One of the most popular optimisers is Adam (Kingma and Ba, 2014). As in SGD, Adam computes the gradients of the loss function with respect to the model's parameters. However, it keeps track of the exponentially decaying averages of past gradients and square gradients (the so called first and second moment, respectively) according to (Kingma and Ba, 2014):

$$\boldsymbol{m}_t = \beta_1 \boldsymbol{m}_{t-1} + (1 - \beta_1) \boldsymbol{\nabla}_w \mathcal{L}(\boldsymbol{w}_{t-1}), \tag{2.21}$$
$$\boldsymbol{v}_t = \beta_2 \boldsymbol{v}_{t-1} + (1 - \beta_2) \boldsymbol{\nabla}_w \mathcal{L}(\boldsymbol{w}_{t-1}). \tag{2.22}$$

The first moment $\boldsymbol{m}_t$ gives an average estimate of the magnitude of the gradients. The second moment $\boldsymbol{v}_t$ instead represents an uncentered variance of the gradients and is a measure of how much the gradients oscillate around a 0 mean gradient. The rates β_1 and β_2 that regulate the exponential decay are usually set to around 0.9. When the training starts, these moments are initialised to zero, causing them to be biased towards zero in the early training stages. To mitigate this bias, Adam applies bias corrections during the first few iterations. This involves adjusting the estimates of the first and second moments by dividing them by correction factors which are different for the mean and the variance:

$$\hat{\boldsymbol{m}}_t = \frac{\boldsymbol{m}_t}{1 - \beta_1^t}, \tag{2.23}$$

$$\hat{\boldsymbol{v}}_t = \frac{\boldsymbol{v}_t}{1 - \beta_2^t}. \tag{2.24}$$

Note the training iteration t at the exponent that makes the bias correction vanish as the training proceeds. Adam then adapts the learning rate separately for each
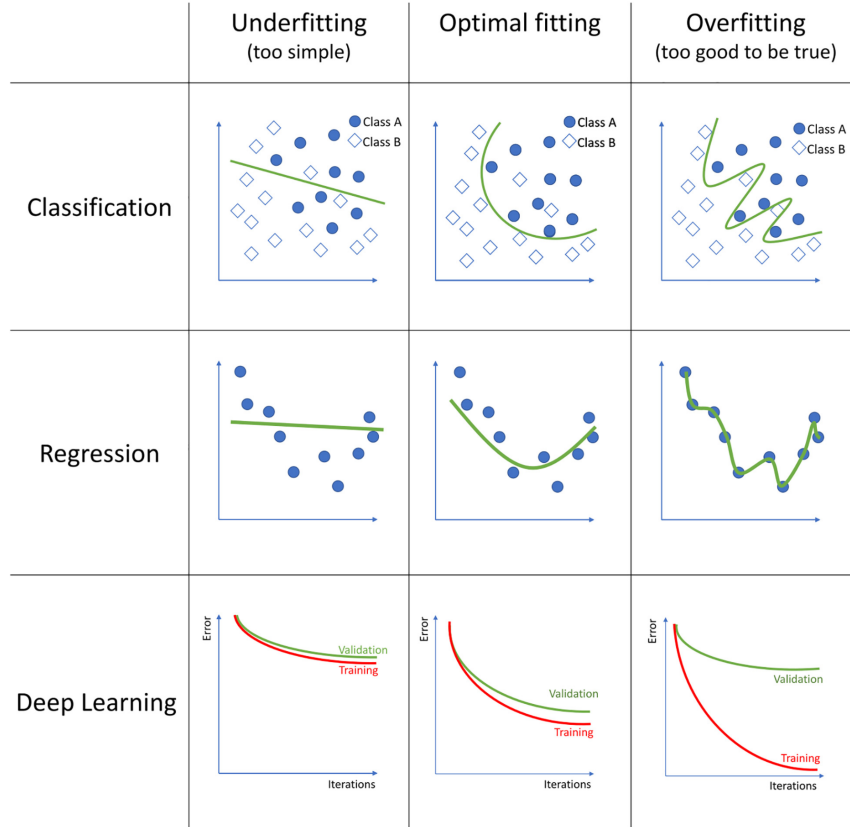
FIGURE 2.6: Examples of model underfitting, model fitting good, and a model overfitting to the data on classification, regression, and deep learning models. In classification, squares and circles represent two different classes. In deep learning, the figures represent the training and validation error over each iteration during the model training process. Figure taken from Solanes and Radua (2022).

weight. This means that weights with a history of large gradients are assigned smaller learning rates, and weights with a history of smaller gradients are assigned larger learning rates. The weight update rule is then given by:

$$w_t = w_{t-1} + \delta w_t = w_{t-1} - \eta \frac{\hat{m}_t}{\sqrt{\hat{v}_t} + \epsilon}, \tag{2.25}$$

where ϵ is a small factor to avoid dividing by zero. This optimisation algorithm helps converge faster and prevents overshooting near minima of the loss landscape. Furthermore, it helps push the optimiser towards the right direction by reducing oscillations during updates. However in general it is not guaranteed that the optimiser converges to a global minimum. If the loss landscape is complex and very irregular it can happen that the optimiser gets trapped in a local minimum (see Figure 2.5).

To summarise, the update reduces the loss and moves the network's predictions closer to the expected output.

2.3.5   Network optimisation and model evaluation

Once the first update of the weights has been performed, the network proceeds to learn by receiving another batch of samples and updating the weights again. The

forward pass, backpropagation, and weights update steps are repeated in cycles until the entire training dataset has been passed in batches through the network. At this point, a single training epoch has been completed. The validation dataset is then used to evaluate the performance of the network on data that it has not seen during the first round of optimisation. The accuracy of the network predictions over the validation dataset is monitored by computing the loss function on the validation dataset. This helps to assess how well the network performs on unseen data and indicates if the training data is overfitted. Signs of overfitting are observed by comparing the evolution of the training and validation losses. If the training and validation loss curves start to diverge during training, with the training loss reducing while the validation loss increases or remains high, this indicates overfitting (see Figure 2.6). A way to mitigate overfitting is to introduce early-stopping during the training which consists of stopping the training process if the validation loss ceases to improve.

The steps outlined above are repeated for a fixed number of iterations (epochs) or until the network achieves a desired level of accuracy on the validation dataset. During each iteration, the forward and backward passes help to refine the network's weights, minimising the loss and improving the network's performance.

After the training is completed, the final weights are saved and the neural network's performance is evaluated on a separate test dataset to gauge its generalisation ability on completely unseen data. While the preparation of the training dataset and the training process itself can be time consuming, once the network has been trained the evaluation on a new set of data is very fast. This is an advantage for example in automatic detection and classification of transient events (see Section 2.6 for some example of applications in astrophysics).

2.4 Simulation based (likelihood-free) inference

In physics and astrophysics, we commonly employ stochastic simulators that aim to emulate some real-world observations or phenomena. Such a simulator incorporates variability and randomness by sampling some variables from probability distributions using Monte Carlo techniques. It relies on theoretical parametric models that depend on certain input parameters $\boldsymbol{\theta} = \{\theta_1, \theta_2, ...\}$. Given this set of input parameters $\boldsymbol{\theta}$, the simulator generates a synthetic realisation of the observed data $\boldsymbol{x}$. In general, underlying the usage of simulation models is the assumption that these models are correct and hence provide an accurate emulation of the real-world processes that generate the observed data. A key challenge in simulation based problems is constraining the free parameters in our simulation models with observational data. The goal is to find the region in this parameter space that is consistent with both the available empirical data and our prior knowledge. In general, we encode our prior knowledge of the phenomenon under study in a prior probability distribution, $\pi(\boldsymbol{\theta})$, over the model parameters $\boldsymbol{\theta}$. The stochastic simulator defines the likelihood probability distribution $\mathcal{P}(\boldsymbol{x}|\boldsymbol{\theta})$ of some data $\boldsymbol{x}$ (that is generated by the simulator or the product of real observations) given some set of model parameters $\boldsymbol{\theta}$. The aim is to compute a posterior probability distribution $\mathcal{P}(\boldsymbol{\theta}|\boldsymbol{x})$ of the model parameters $\boldsymbol{\theta}$ given the data $\boldsymbol{x}$. This can be achieved by using the Bayes theorem (Stuart and Ord, 1994):

$$\mathcal{P}(\boldsymbol{\theta}|\boldsymbol{x}) = \frac{\pi(\boldsymbol{\theta})\mathcal{P}(\boldsymbol{x}|\boldsymbol{\theta})}{\mathcal{P}(\boldsymbol{x})} \tag{2.26}$$
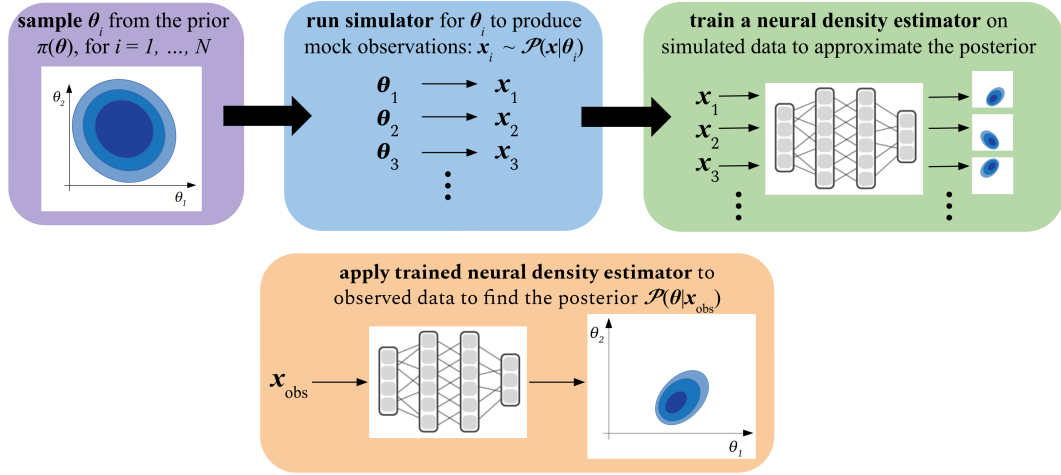
FIGURE 2.7: Neural posterior estimation workflow
Scheme of the workflow for the neural posterior estimation (NPE) algorithm (See text for more details).

where $\mathcal{P}(\boldsymbol{x}|\boldsymbol{\theta})$ is the likelihood of our data, $\boldsymbol{x}$, given the parameter, $\boldsymbol{\theta}$, and

$$\mathcal{P}(\boldsymbol{x}) \equiv \int \mathcal{P}(\boldsymbol{x}|\boldsymbol{\theta}')\pi(\boldsymbol{\theta}')\,\mathrm{d}\boldsymbol{\theta}', \tag{2.27}$$

denotes the evidence obtained by marginalizing over all $\boldsymbol{\theta}$. where $\mathcal{P}(\boldsymbol{x})$. However, for complex simulators, the likelihood $\mathcal{P}(\boldsymbol{x}|\boldsymbol{\theta})$ is typically intractable as it is implicitly defined by the simulator. Even if it were tractable, for multidimensional problems the posterior computation can become very costly as Equation (2.27) involves an integral over $\boldsymbol{\theta}$, which becomes challenging for simulators with high-dimensional parameter spaces. Simulation-based likelihood-free inference aims to bypass these problems and directly compute the posterior distribution $\mathcal{P}(\boldsymbol{\theta}|\boldsymbol{x})$ without having access to the likelihood probability density (see Cranmer et al., 2020, for a review).

Deep neural networks are particularly useful for this task because they can be used to learn a (probabilistic) association between data and the underlying model parameters, i.e., extract a posterior probability distribution from simulated data without the need to explicitly compute the likelihood. There are three main approaches to achieve this goal:

- **Neural posterior estimation (NPE)**: The network learns to directly map the simulator output, $\boldsymbol{x}$, onto the posterior distribution, $\mathcal{P}(\boldsymbol{\theta}|\boldsymbol{x})$, for the underlying parameters, $\boldsymbol{\theta}$. This requires the use of a flexible neural density estimator such as a normalising flow or a mixture density network (MDN) (e.g. Papamakarios and Murray, 2016; Lueckmann et al., 2017; Greenberg, Nonnenmacher, and Macke, 2019; Mishra-Sharma and Cranmer, 2022; Vasist et al., 2023; Dax et al., 2021; Hahn et al., 2023).

- **Neural likelihood estimation (NLE)**: The network emulates the simulator by learning an association between $\boldsymbol{\theta}$ and $\boldsymbol{x}$, thus providing direct access to an approximation of the likelihood, $\mathcal{P}(\boldsymbol{x}|\boldsymbol{\theta})$. Because the prior is known, the posterior can then be obtained by an additional Markov chain Monte Carlo (MCMC) sampling step (e.g., Papamakarios, Sterratt, and Murray, 2018; Alsing et al., 2019).

- **Neural ratio estimation (NRE)**: Here, the network learns the likelihood-to-evidence ratio, $r(\theta, x) \equiv \mathcal{P}(x|\theta)/\mathcal{P}(x)$, which is equivalent to $\mathcal{P}(\theta|x)/\mathcal{P}(\theta)$ using Bayes' theorem (6.25). Once $r(\theta, x)$ is known, the posterior can be recovered through MCMC by sampling the prior weighted by the ratio, $r(\theta, x)$ (e.g., Hermans, Begy, and Louppe, 2019; Miller et al., 2021; Bhardwaj et al., 2023).

In the following I will concentrate on NPE as it allows to directly derive an estimation of the posterior distribution without the extra step of a MCMC sampling. As we will adopt the NPE method in Chapter 6, in this section I will discuss further details of this approach.

2.4.1 Neural posterior estimation

Typically, in the NPE approach, a neural network, F, is trained to map a given simulation outcome x onto the parameters ψ of a density estimator q_ψ which is used to approximate the posterior distribution, $\mathcal{P}(\theta|x)$, of the underlying model parameters θ (Papamakarios and Murray, 2016; Lueckmann et al., 2017; Greenberg, Nonnenmacher, and Macke, 2019, see also the schematic workflow in Figure 2.7). To train the network, suppose that we have a dataset $\{(\theta_i, x_i)\}$ with $i = 1, ..., N$, where $\{\theta_i\}$ is a set of model parameters drawn from the prior distribution $\pi(\theta)$ and $\{x_i\}$ is the set of corresponding simulation outcomes. For a given sample (θ_i, x_i), we define the following loss function:

$$\mathcal{L}(\theta_i) = -\log q_{F(x_i, w)}(\theta_i), \tag{2.28}$$

where we make explicit that the density estimator parameters depend on the input x_i and the network weights w, i.e., $\psi = F(x_i, w)$. Minimising this loss function means adjusting the network weights so that the probability defined by the density estimator of the label parameters θ_i is maximised. As demonstrated in Papamakarios and Murray (2016) for sufficiently complex networks and flexible density estimators, minimising the expectation value of the loss in Equation (2.28) leads to $q_{F(x, w)}(\theta) \approx \mathcal{P}(\theta|x)$ as $N \to \infty$. This can be seen by using the Kullback-Leibler divergence, $D_{KL}(\mathcal{P}_1 || \mathcal{P}_2)$, which is a measure of the difference between two probability distributions, $\mathcal{P}_1$ and $\mathcal{P}_2$ (Kullback and Leibler, 1951). First we note that $q_{F(x, w)}(\theta) \approx \mathcal{P}(\theta|x)$ is equivalent to $\mathcal{P}(x|\theta)\pi(\theta) \approx \mathcal{P}(x)q_{F(x, w)}(\theta)$ by using Bayes theorem (Equation (2.26)), so that we can write the following (Papamakarios and Murray, 2016):

$$
\begin{aligned}
D_{KL}[\mathcal{P}(\theta|x) \,||\, q_{F(x, w)}(\theta)] &= D_{KL}[\mathcal{P}(x|\theta)\pi(\theta) \,||\, \mathcal{P}(x)q_{F(x, w)}(\theta)] \\
&\equiv \int \mathcal{P}(x|\theta)\pi(\theta) \log \frac{\mathcal{P}(x|\theta)\pi(\theta)}{\mathcal{P}(x)q_{F(x, w)}(\theta)} d\theta dx \\
&= \int \mathcal{P}(x|\theta)\pi(\theta) \log \frac{\mathcal{P}(\theta|x)}{q_{F(x, w)}(\theta)} d\theta dx \\
&= \int \mathcal{P}(x|\theta)\pi(\theta) \log \mathcal{P}(\theta|x) d\theta dx \\
&\quad - \int \mathcal{P}(x|\theta)\pi(\theta) \log q_{F(x, w)}(\theta) d\theta dx \\
&= \mathrm{E}_{\mathcal{P}(x|\theta)\pi(\theta)}[\log \mathcal{P}(\theta|x)] - \mathrm{E}_{\mathcal{P}(x|\theta)\pi(\theta)}[\log q_{F(x, w)}(\theta)].
\end{aligned}
\tag{2.29}
$$

The expectation values are computed over the joint probability distribution $\mathcal{P}(x, \theta) = \mathcal{P}(x|\theta)\pi(\theta)$ that can be sampled using the simulator. The first expectation value does not depend on the network parameters. Therefore it is a constant term that can be neglected in the minimisation procedure. This last expression shows that minimising the Kullback-Leibler divergence is equivalent to minimising the expectation value of the loss function defined by Equation (2.28).

A commonly employed density estimator is masked autoregressive flow (MAF). A MAF is a composition of invertible and differentiable transformations that are applied to a simple distribution, such as a standard Gaussian distribution, to transform it into more complex distributions (Jimenez Rezende and Mohamed, 2015; Papamakarios, Pavlakou, and Murray, 2017). Each transformation is parametrised by neural networks which depend on the input data x. In this way, MAFs can be trained to approximate any target distribution.

An alternative to MAFs are mixture density networks (MDNs). I will focus on this type of architecture as we will employ it in Chapater 6. MDNs can model complex probability distributions by approximating them using a mixture of distributions. A mixture of distributions (also known as a mixture model) is a statistical model that represents the probability distribution of a random variable as a combination of two or more component distributions. A MDN thus takes some features x as input and outputs the parameters ψ defining the mixture model. More precisely, these include the parameters that determine the shape and characteristics of each component distribution and the mixing coefficients that determine their weights in the mixture model. An example of a mixture model is a mixture of multivariate Gaussian distributions in the D-dimensional space that is defined as:

$$q_{\psi}(\theta) = \sum_{c=1}^{C} \alpha_c \mathcal{N}(\theta|\mu_c, \Sigma_c), \tag{2.30}$$

where C is total the number of Gaussian components $\mathcal{N}$, α_c are the mixing coefficients, μ_c is the vector of means and Σ_c is the covariance matrix of the c Gaussian component defined as:

$$\mathcal{N}(\theta|\mu_c, \Sigma_c) = \frac{1}{(2\pi)^{n/2} \det(\Sigma_c)^{1/2}} \exp\left[-\frac{1}{2}(\theta - \mu_c)^{\mathrm{T}} \Sigma_c^{-1} (\theta - \mu_c)\right]. \tag{2.31}$$

In this case, the mixture density network outputs the parameters α_c, μ_c and the entries of the covariance matrix Σ_c for every Gaussian component. To evaluate the loss for a given simulated data sample (θ_i, x_i), we need to compute:

$$\mathcal{L}(\theta_i) = -\log \sum_{c=1}^{C} \alpha_c \mathcal{N}(\theta_i|\mu_c, \Sigma_c), \tag{2.32}$$

where the parameters α_c, μ_c and Σ_c depend on the input features x_i and the network weights w. Since we have an exponential function in the normal distribution, we might obtain very small numbers when evaluating the loss in this form. Such values might be too small to be represented with the correct precision by the computing system (numerical underflow). To avoid this issue, we take advantage of the fact

that $x = \exp(\log x)$. The loss function can thus be rewritten in the following form:

$$
\begin{aligned}
\mathcal{L}(\boldsymbol{\theta}) &= -\log \sum_{c=1}^{C} \exp\left[\log\left(\alpha_c \mathcal{N}(\boldsymbol{\theta}|\boldsymbol{\mu}_c, \boldsymbol{\Sigma}_c)\right)\right] \\
&= -\log \sum_{c=1}^{C} \exp\left[\log \alpha_c - \frac{n}{2}\log(2\pi) - \frac{1}{2}\log\det(\boldsymbol{\Sigma}_c) - \frac{1}{2}(\boldsymbol{\theta} - \boldsymbol{\mu}_c)^{\mathrm{T}} \boldsymbol{\Sigma}_c^{-1} (\boldsymbol{\theta} - \boldsymbol{\mu}_c)\right].
\end{aligned}
$$
(2.33)

Implementing this log of the sum of exponentials numerically allows a stable calculation of the loss.

Once trained on a large collection of simulations, this type of network is able to generate a precise estimation of the posterior density distribution for any given input of empirical data. This allows a fast inference procedure which is also amortised, meaning that it does not rely on a single specific observation.

2.5   Diagnostic checks

After a density estimator has been trained on a dataset of simulations and to produce posteriors, its predictions should undergo several diagnostic tests, before being used for inference on actual observed data. A first qualitative check is the so-called posterior-predictive check. It consists of inspecting if the data generated from the simulator using the parameters sampled from the posterior looks similar to the actual input data used to estimate the posterior. Even if not fully systematic, this test gives already an indication if the predicted posterior is biased or inaccurate. In the following, I report two more quantitative tests that should be performed to asses if the trained estimator is well calibrated (see Chapter 6 for an application of these tests).

2.5.1   Simulation based calibration

Talts et al. (2018) proposed a method called simulation-based calibration (SBC) that provides a qualitative view and a quantitive measure to check if posterior predictions are trustworthy. SBC can diagnose whether the estimated posterior predictions are biased, i.e., systematically overestimating or underestimating the true parameter values, and whether the posterior uncertainties are unbalanced, i.e., either over-confident or under-confident. Therefore, SBC can be viewed as a necessary (but not sufficient) condition for a valid inference algorithm. If SBC checks fail, the inference results should not be trusted. On the other hand, if SBC checks pass, the posterior estimation should be well calibrated but this is no guarantee for the method to be working properly.

SBC is based on ranking parameters $\boldsymbol{\theta}$ sampled from the prior $\pi(\boldsymbol{\theta})$ and checking their rank statistics (see also Cook, Gelman, and Rubin, 2006). The procedure is as follows. We consider a test dataset $\{(\boldsymbol{\theta}_i, \boldsymbol{x}_i)\}$ (that has been used neither for training nor validating the network) and compute the posterior $\mathcal{P}(\boldsymbol{\theta}|\boldsymbol{x}_i)$ by processing each $\boldsymbol{x}_i$ through the trained network. From each posterior $\mathcal{P}(\boldsymbol{\theta}|\boldsymbol{x}_i)$, we sample S values $\boldsymbol{\theta}_s$ and count how many of these samples are smaller than the original $\boldsymbol{\theta}_i$. This count gives the rank associated with the parameter $\boldsymbol{\theta}_i$. Once this has been performed for the entire dataset $\{(\boldsymbol{\theta}_i, \boldsymbol{x}_i)\}$, we analyze the obtained rank statistics. For a well-calibrated posterior, the ranks follow a uniform distribution (Talts et al., 2018). If this is not the case, we obtain an indication that the estimated posteriors are not
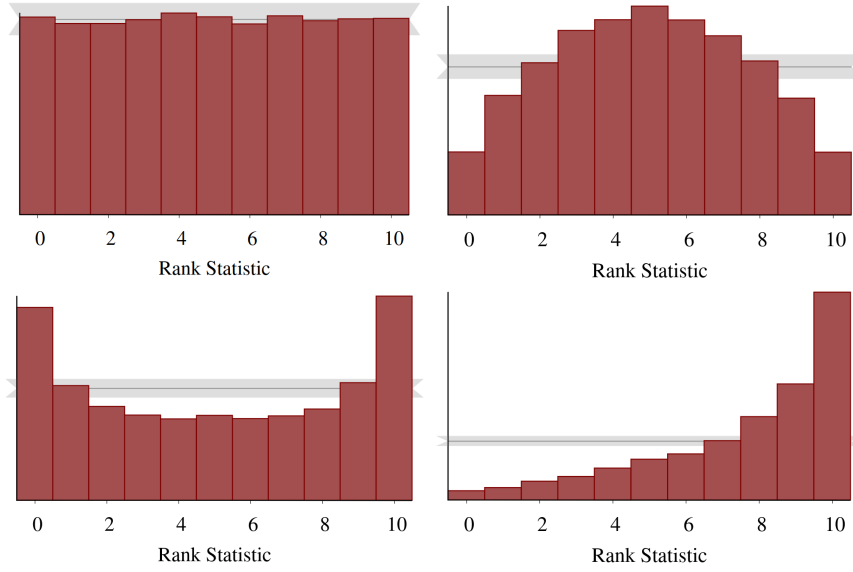
FIGURE 2.8: Example sketches showing the distribution of the ranks statistics in the case of a density estimator predicting well calibrated posteriors (*top left plot*), posteriors that are too conservative (*top right plot*), overconfident posteriors (*bottom left plot*) and posteriors that are biased, i.e., that underestimate the parameter values in this specific case (Figures taken from Talts et al., 2018).

accurate on average (see Figure 2.8). For example, if the rank distribution is skewed to the left, i.e., it is dominated by lower ranks, the predicted posteriors tend to overestimate the parameters θ. If the rank distribution is skewed to the right, i.e., it is dominated by high ranks, the predicted posteriors tend to underestimate the parameters θ. Moreover, if the rank distribution has a "U" shape, and is dominated by low and high ranks, the predicted posteriors tend to be overconfident, that is they underestimate the uncertainties on the parameters θ. Finally, if the rank distribution has a "∩" shape and is dominated by medium ranks, the predicted posteriors tend to be too conservative. This implies that they are overestimating the uncertainties on the parameters θ. As θ is usually an array of parameters, the rank statistics are computed over the marginalised 1D posterior distributions of each these parameters. Therefore, this test might not be particularly sensitive to the shape of the posterior in the multi-dimensional space and the presence of correlations between parameters. While rank statistics are a useful metric to assessing the quality of a posterior distribution, we will primarily focus on another method outlined below in the remainder of this thesis.

2.5.2 Coverage probability

Another test to asses if a posterior estimator is well calibrated is checking its coverage probability (Hermans et al., 2021). As in SBC, we consider a test dataset $\{(\theta_i, x_i)\}$ and compute the posterior $\mathcal{P}(\theta|x_i)$ for each x_i. For each posterior, we then define a credible region Θ_i with the smallest volume in the multidimensional parameter space of θ corresponding to a total probability $1 - \alpha$ with $\alpha \in [0, 1]$:

$$\int_{\Theta_i} \mathcal{P}(\theta|x_i)d\theta = 1 - \alpha, \tag{2.34}$$
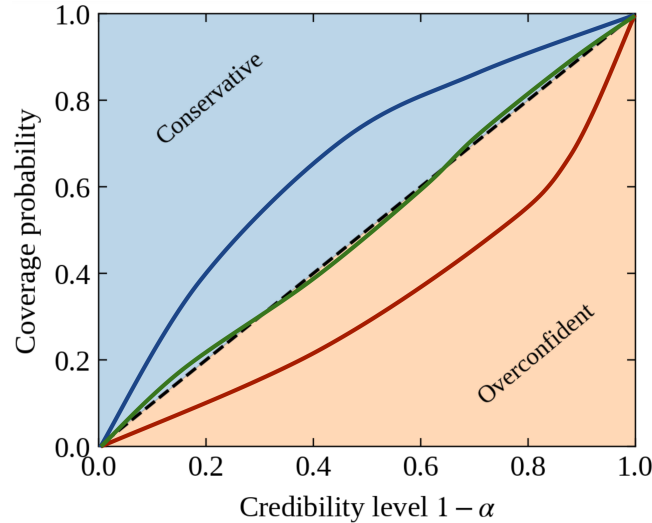
FIGURE 2.9: Coverage probability plot showing the computed coverage probability versus the credibility level $1 - \alpha$ for overconfident posterior distributions (*red curve*), well-calibrated posteriors (*green curve*) and overconfident posteriors (*blue curve*).

where $1 - \alpha$ defines the so-called credibility level. Considering the region with the smallest volume guarantees that we are focusing on the region of the parameter space that encloses the parameter values $\boldsymbol{\theta}$ with the highest posterior probability density. In the literature, this is also called the highest posterior density region (Hyndman, 1996).

By counting how many $\boldsymbol{\theta}_i$ fall inside the corresponding credibility regions Θ_i we obtain a measure of how well the estimated posteriors are able to recover the label parameters $\boldsymbol{\theta}$. This count gives an estimate of the so-called coverage probability. If the posterior estimator is well calibrated, the coverage probability should be equal to the value $1 - \alpha$. If the coverage probability is higher than $1 - \alpha$, this is a symptom for a posterior estimator that tends to generate posteriors that are too conservative. On the other hand, if the coverage probability is lower than $1 - \alpha$, this indicates that the estimated posteriors are overconfident. We will use this test in Chapter 6 (see also Appendix D for more details and Figure 2.9 for an illustrative representation).

2.6 Applications in astrophysics

In recent years there has been an increasing number of works that apply the machine learning techniques outlined above in the field of astronomy and astrophysics. In particular, for neutron-star and compact-object related science, ML algorithms have for example been developed to classify new pulsars candidates (see for example Bethapudi and Desai, 2018; Balakrishnan et al., 2021; Lin, Li, and Luo, 2020) as well as transient radio events such as fast radio bursts (FRBs) (Connor and van Leeuwen, 2018; Agarwal et al., 2020) or to visualize from another perspective the pulsar population through minimum spanning trees (García, Torres, and Patruno, 2022). Other approaches have aimed to combine classification and regression tasks to forecast and analyse gravitational-wave signals in real-time and perform parameter estimation for the merger events (George and Huerta, 2018; Cabero, Mahabal, and McIver, 2020; Wei and Huerta, 2020; Skliris, Norman, and Sutton, 2020; Gerosa, Pratten, and Vecchio, 2020). Some other studies focused on reconstructing the neutron-star

equation of state from observed quantities, like their masses and radii (Fujimoto, Fukushima, and Murase, 2018; Morawski and Bejger, 2020). In Chapter 5 we will study the feasibility of using deep neural networks (in particular CNNs) to infer the dynamical properties of the neutron star population.

In addition, the rapidly developing field of simulation-based inference is now offering new tools to perform parameter estimation with complex simulators (see Cranmer et al., 2020, for a review). For example, SBI has been employed for parameter estimation from gravitational-wave signals (Green, Simpson, and Gair, 2020; Bhardwaj et al., 2023). Vasist et al. (2023) inferred the physical properties of exoplanetary atmospheres from spectroscopic data. Moreover, Khullar et al. (2022) have been able to retrieve galaxy properties from their spectral energy distributions (SEDs). In the context of cosmology, Jeffrey, Alsing, and Lanusse (2021) used an SBI framework to infer cosmological parameters from weak lensing maps. In this work, we are going to employ for the first time an SBI framework developed by Tejero-Cantero et al. (2020) in conjunction with a population synthesis simulator to retrieve the birth properties of the Galactic population of isolated radio pulsars (see Chapter 6).

Chapter 3

Long-period pulsars as possible outcomes of supernova fallback accretion

3.1 Introduction

The spin-period distribution of the pulsar population reflects intrinsic properties of neutron-star formation, early evolution, magnetic-field decay, and age (see sec. 1.13). Until a few years ago, the spin distribution of observed pulsars was ranging between ~ 0.002–12 s. At the fastest extreme, we have recycled millisecond pulsars (mostly in binaries), while the slowest extreme is populated by magnetars (see Figure 1.12). The historical lack of isolated pulsars with periods $\gtrsim 12$ s has been intriguing and interpreted in different ways, ranging from the presence of a death line below which radio emission is quenched (Ruderman and Sutherland, 1975; Bhattacharya and van den Heuvel, 1991; Chen and Ruderman, 1993, see also Section 1.9.3), to magnetic-field decay coupled with the presence of a highly resistive layer in the inner crust (possibly due to the existence of a nuclear pasta phase Pons, Viganò, and Rea, 2013, see also Section 1.3), as well as due to an observational bias caused by high band pass filters in radio searches (albeit not in X-ray searches). Although the main reason is uncertain, all of these effects likely contribute at some level to the observational paucity of long-period pulsars (Wu et al., 2020).

However, recent radio surveys, in particular thanks to new radio interferometers such as the LOw Frequency ARray (LOFAR; van Haarlem et al., 2013), MeerKAT (Jonas, 2009), Australian SKA Pathfinder (ASKAP; Hotan et al., 2021), and the Murchison Widefield Array (MWA; Tingay et al., 2013; Wayth et al., 2018), have started to uncover the existence of a new population of slowly rotating radio pulsars that challenge our understanding of the pulsar population and its evolution.

Two radio pulsars, PSR J1903+0433 (Han et al., 2021) and PSR J0250+5854 (Tan et al., 2018), have been recently discovered with periods of 14 s and 23 s, respectively. Moreover, a ~ 76 s radio pulsar with a magnetic field of $B \sim 1.3 \times 10^{14}$ G (PSR J0901-4046, Caleb et al., 2022) and two peculiar radio transient with a periodicity of ~ 1091 s (GLEAM-X J162759.5-523504.3; Hurley-Walker et al., 2022a) and ~ 1318 s (GPM J1839–10 Hurley-Walker et al., 2023) have been discovered. These latter sources are currently still uncertain in nature. In particular, GLEAM-X J1627 has a very variable flux, showing periods of "radio outburst" lasting a few months, a 90% linear polarisation, and a very spiky and variable pulse profile. GPM J1839–10 is also characterised by a very irregular and variable pulse profile and has been active for at least 34 years. Several interpretations have been advanced to explain the

mysterious nature of these sources. As argued in the discovery papers, their emission characteristics are typical of observed radio magnetars (Kaspi and Beloborodov, 2017; Esposito et al., 2020) (although their magnetic fields are still poorly constrained). Alternatively, they could be strongly magnetised white dwarfs having spun down to the observed long spin period due to their larger moment of inertia (see also Tong (2022) and Chapter 4 for a discussion of the nature of these sources). Furthermore, a few years ago the X-ray emitting neutron star 1E 161348-5055 at the centre of the 2 kyr old supernova remnant (SNR) RCW103, with a measured modulation of ~ 6.67 hr, showed a large magnetar-like X-ray outburst (Rea et al., 2016; D'Aì et al., 2016), demonstrating the source's isolated magnetar nature despite its long period and young age.

In general, studying the possible mechanisms that could lead to the formation of long-period pulsars is of great interest in order to understand their origin, eventual links with other types of neutron stars (see Section 1.13) and possible connections to periodic activity of transient events such as fast radio bursts (see for example Beniamini, Wadiasingh, and Metzger, 2020; Xu et al., 2021). One possible avenue that could lead to enhanced spin-down, especially in magnetars, is the presence of mass-loaded charged particle winds and outflows that could be particularly active after giant-flare episodes and temporarily expand the open magnetic flux region of the star (see for example Thompson et al., 2000; Tong et al., 2013; Beniamini, Wadiasingh, and Metzger, 2020). Another possibility is the interaction of the neutron star with a fallback disk formed after the supernova explosion. Soon after their formation, neutron stars will necessarily witness fallback accretion with different accretion rates depending on the progenitor properties and explosion dynamics (Ugliano et al., 2012; Perna et al., 2014; Janka, Wongwathanarat, and Kramer, 2022). If the fallback matter possesses sufficient angular momentum, it could form a long-lasting accretion disk that will interact with the neutron star. For certain ranges of the initial spin period P_0, magnetic field B_0 and disk accretion rate $\dot{M}_{\rm d,0}$, fallback after the supernova explosion can substantially affect the pulsar spin evolution, in some cases slowing down the pulsar period significantly more than standard dipolar spin-down losses alone. In this context, the long period observed in 1E 161348-5055 has been interpreted as the spin period of a magnetar interacting with a fallback disk in a propeller state by numerous authors (Li, 2007; Ho and Andersson, 2017; Tong et al., 2016; Xu and Li, 2019). Moreover Chatterjee, Hernquist, and Narayan (2000), Ertan et al. (2009), and Benli and Ertan (2016) developed detailed numerical models of the interaction between a fallback disk and a neutron star to explain the emission of magnetars. Although appealing, this model struggles to explain several observed features of these types of sources, such as their burst and flaring activity (see Mereghetti, 2013; Turolla, Zane, and Watts, 2015; Kaspi and Beloborodov, 2017; Esposito, Rea, and Israel, 2021, for a review). Overall, these studies on fallback disk accretion suggest that the presence or absence of a fallback disk around newly born neutron stars could be a determining factor for their evolution in the P-$\dot{P}$ diagram and (Figure 1.12) as well as their emission properties and could further be relevant to explain the connections between different neutron star classes (see for example Alpar, Ankay, and Yazgan, 2001, and Section 1.13).

In this chapter based on the work published in Ronchi, Rea, Graber, and Hurley-Walker (2022), we revisit the fallback scenario to specifically analyze the spin-period evolution of newly born pulsars that witness accretion from a fallback disk (Section 3.3), and determine the parameter ranges that allow pulsars to experience efficient spin-down in the presence of magnetic field decay. Furthermore, we study the characteristics of the newly discovered long-period radio sources GLEAM-X J1627
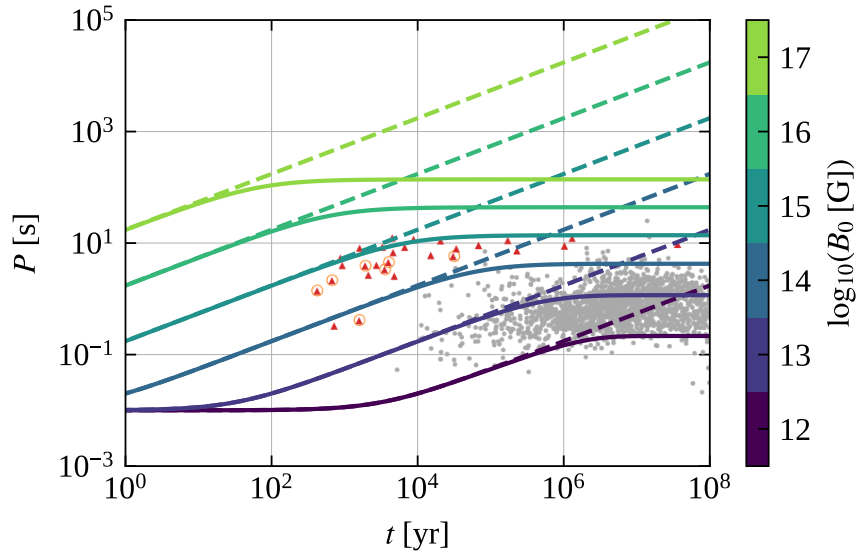
FIGURE 3.1: The evolution of the pulsar spin-period P in time t assuming dipolar spin-down for different values of the initial magnetic field B_0 (indicated by the colour code in the colour bar). The dashed lines represent the evolution curves for a constant magnetic field, while the solid lines represent the evolution tracks for a decaying magnetic field according to Equation (3.4). In the background, the gray points represent the observed radio pulsars (data from the ATNF Pulsar Catalog, https://www.atnf.csiro.au/research/pulsar/psrcat/; Manchester et al., 2005) and red triangles represent the currently detected magnetars (data from the Magnetar Outburst Online Catalogue, http://magnetars.ice.csic.es/; Coti Zelati et al., 2018) highlighting those that exhibit radio emission (orange circles). We consider the characteristic age $\tau_c = P/(2\dot{P})$ as a proxy for their real age.

(see Section 3.6), GPM J1839–10 (see Section 3.7) and pulsar PSR J0901-4046 (see Section 3.8) in the context of the fallback scenario in order to constrain their nature and evolution (see Section 3.9). We provide a summary in Section 3.10. [1]

3.2 Period evolution of pulsars slowing down via dipolar losses

Ordinary rotation-powered pulsars are expected to slow down via electromagnetic dipolar losses, with an additional component driven by their magnetic-field decay (see Section 1.8). In this scenario, if $\omega = 2\pi/P$ is the angular spin frequency of the pulsar, the electromagnetic torque causes the star to slow down according to (see Equation (1.40)):

$$\frac{d\omega}{dt} = -\beta B(t)^2 \omega^3, \tag{3.1}$$

where for an aligned rotator $\beta \sim R_{NS}^6/(4c^3 I_{NS}) \sim 1.5 \times 10^{-41}\,\mathrm{s\,G^{-2}}$ assuming a typical neutron-star radius $R_{NS} \sim 11\,\mathrm{km}$, and moment of inertia $I_{NS} \sim 1.5 \times 10^{45}\,\mathrm{g\,cm^2}$. As discussed in Section 1.8 the inclination angle dependence and uncertainty on the

---

[1]Jupyter Notebooks to reproduce the plots and results in this chapter are publicly available at https://github.com/MicheleRonchi/pulsar_fallback

neutron-star mass and radius introduce a correction of at most an order of magnitude to the value of β in the equation above (see also Spitkovsky, 2006).

If we consider a constant magnetic field $B(t) = B_0$, we can solve Equation (3.1) to find the spin-period evolution in time:

$$P(t) = P_0 \left(1 + \frac{t}{t_{em}}\right)^{1/2}, \tag{3.2}$$

where we define an electromagnetic spin-down timescale $t_{em} = P_0^2 / \left(8\pi^2 \beta B_0^2\right)$. Note that for $t \gg t_{em}$ the period evolution loses memory of the initial spin period, P_0, and only depends on the magnetic field value:

$$P(t) = \left(8\pi^2 \beta B_0^2 t\right)^{1/2}. \tag{3.3}$$

This means that neutron stars born with different initial spin periods but same magnetic field strength will end up evolving identically.

However, several studies have shown that neutron star magnetic fields evolve over time and decay due to the combined action of Ohmic dissipation and the Hall effect in the star's crust (Pons and Geppert, 2007; Pons, Miralles, and Geppert, 2009; Viganò et al., 2013; Pons and Viganò, 2019; De Grandis et al., 2020, see also Section 1.6). For simplicity, if we consider a crustal-confined magnetic field, we can use the phenomenological description of the field decay presented in Aguilera, Pons, and Miralles (2008a) and Aguilera, Pons, and Miralles (2008b), given by the analytic expression (see also Equation (1.15)):

$$B(t) = B_0 \frac{e^{-t/\tau_{Ohm}}}{1 + \frac{\tau_{Ohm}}{\tau_{Hall,0}} \left[1 - e^{-t/\tau_{Ohm}}\right]}. \tag{3.4}$$

This equation captures a first stage that is characterised by rapid (non-exponential) decay and regulated by the Hall timescale $\tau_{Hall,0}$, and a second stage that is characterised by exponential decay due to Ohmic dissipation and regulated by the timescale τ_{Ohm}. These two characteristic timescales were defined in eqs. (1.13) and (1.14). I report them here again as easy reference:

$$\tau_{Hall,0} = \frac{4\pi e n_e L^2}{c B_0}$$
$$\simeq 6.4 \times 10^4 \, \text{yr} \left(\frac{n_e}{10^{35} \, \text{cm}^{-3}}\right) \left(\frac{L}{1 \, \text{km}}\right)^2 \left(\frac{B_0}{10^{14} \, \text{G}}\right)^{-1}, \tag{3.5}$$

$$\tau_{Ohm} = \frac{4\pi \sigma L^2}{c^2}$$
$$\simeq 4.4 \times 10^6 \, \text{yr} \left(\frac{\sigma}{10^{24} \, \text{s}^{-1}}\right) \left(\frac{L}{1 \, \text{km}}\right)^2, \tag{3.6}$$

where we remind that σ is the dominant conductivity based on phonon or impurity scattering and L is the typical lengthscale over which the relevant physical quantities (i.e., n_e, B_0 and σ) change inside the crust (see Cumming, Arras, and Zweibel, 2004; Gourgouliatos and Cumming, 2014a). Note also that in Equation (3.4), $\tau_{Hall,0}$ represents the Hall timescale for the initial magnetic field strength B_0 (we refer to Aguilera, Pons, and Miralles, 2008a, for more details). In general, the value of the Hall and the Ohmic timescales can vary by orders of magnitude within the crust and during the evolution, depending strongly on the density profile and magnetic

field intensity and curvature. The Ohmic timescale additionally varies with temperature due to its dependence on the conductivity (see Section 1.6). Although a rough simplification, Equation (3.4) and the timescales defined above are consistent with the evolution history inferred for Galactic magnetars (Colpi, Geppert, and Page, 2000; Dall'Osso, Granot, and Piran, 2012; Beniamini et al., 2019) and are able to qualitatively capture the magnetic field evolution obtained with more sophisticated magneto-thermal numerical simulations (see for example Viganò et al., 2013; De Grandis et al., 2020; Viganò et al., 2021).

Using these prescriptions for the magnetic field, we solve Equation (3.1) and determine the evolution of the spin-period in time. In Figure 3.1, we present the corresponding behaviour for B_0 values in the range 10^{12-17} G. Since the long-term evolution for $t \gg t_{\rm em}$ is insensitive to the initial spin period value, we fix P_0 to a fiducial value of 10 ms which is compatible with the birth spin-period distributions inferred from population synthesis studies (see for example Faucher-Giguère and Kaspi, 2006; Gullón et al., 2014; Cieślar, Bulik, and Osłowski, 2020). After an initial phase of duration $t_{\rm em}$, where the spin period remains almost constant and equal to its initial value, P starts to evolve $\propto t^{1/2}$. If the magnetic field remains constant, this evolution proceeds indefinitely (dashed lines). If B decays over time according to Equation (3.4), the electromagnetic torque eventually becomes negligible and the spin period stops increasing and stabilises (solid lines).

From this plot, it is evident that, if one assumes a decaying magnetic field in the neutron star crust, explaining the existence of slowly rotating neutron stars with spin periods $P \gtrsim 100$ s becomes problematic since one would require rather extreme magnetic field values ($B \gtrsim 10^{16}$ G). If instead, some mechanism prevents the crustal magnetic field from decaying in time, the star could reach longer spin periods more easily since a strong electromagnetic torque is maintained over longer times. After a few Hall timescales, the magnetic field in the crust reorganises towards smaller scales through the Hall cascade (e.g. Brandenburg, 2020) and approaches a quasi-equilibrium configuration during which the dissipation of the magnetic field is slowed down. In this phase, commonly named the "Hall attractor" (Gourgouliatos and Cumming, 2014b), the electric currents are predominantly confined to the inner crust where subsequent dissipation occurs on the Ohmic timescale, which in turn depends on the local properties of the inner crust. As a result, magnetic field decay cannot be halted indefinitely but will proceed on timescales of a few Myr (see however Pons, Viganò, and Rea, 2013, on how the pasta phase could greatly shorten these timescales). Another possibility for maintaining strong magnetic fields is that the electric currents could be predominantly present in the neutron star core allowing the magnetic field to stay stable and barely decay over time (see for example Viganò et al., 2021). However it is still unclear if such conditions can be realised since little is known about magnetic field evolution in neutron star cores (in particular the effects of superfluid and superconducting components on field evolution is uncertain; Graber et al., 2015; Passamonti et al., 2017; Ofengeim and Gusakov, 2018; Gusakov, Kantor, and Ofengeim, 2020). In general, we cannot exclude the possibility that long-period pulsars are the result only of electromagnetic spin-down in the presence of strong and persistent magnetic fields, possibly supported by a Hall attractor or a core component. However, in light of our current knowledge of magnetic field evolution in neutron stars, we suggest that dipolar spin-down alone struggles to explain the existence of long-period pulsars. Instead, we argue below that these sources could be the result of a different scenario, whose ingredients are readily available in standard neutron-star formation models.

3.3 Accretion from a fallback disk

An alternative scenario that could explain the existence of strongly magnetised and slowly rotating pulsars involves the interaction between a newly born highly magnetic neutron star and fallback material from the supernova explosion.

In the early stages after the supernova explosion, the proto-neutron star emits a powerful neutrino wind which exerts a pressure on the outer envelope of the progenitor star. The duration of this wind is believed to be of the order of $\sim 10\,\text{s}$ after core bounce, which corresponds to the neutrino-cooling timescale for a newborn neutron star (Ugliano et al., 2012; Ertl et al., 2016). On this timescale, the magnetosphere of the neutron star reaches an equilibrium configuration in the region swept by the neutrino wind. After the subsiding of this neutrino-driven wind, the gas in the inner envelopes of the exploding star decelerates due to the gravitational pull of the neutron star and collapses back. This leads to the onset of fallback accretion. According to simulations (Ugliano et al., 2012; Ertl et al., 2016; Janka, Wongwathanarat, and Kramer, 2022), the total fallback mass can reach values up to $M_{\text{fb}} \lesssim 0.1 M_\odot$, while the fallback mass rate can reach values of around $\sim 10^{27-31}\,\text{g s}^{-1}$ in the first $\sim 10 - 100\,\text{s}$ after bounce, afterwards decreasing according to a power law $\sim t^{-5/3}$, which is compatible with theoretical predictions for spherical supernova fallback (Michel, 1988; Chevalier, 1989). However, if part of the fallback matter possesses sufficient angular momentum, it will circularise to form an accretion disk.

Mineshige et al. (1997) and Menou, Perna, and Hernquist (2001) have studied the formation and time evolution of fallback accretion disks around compact objects. In particular, by using the self-similar solution of Cannizzo, Lee, and Goodman (1990), Menou, Perna, and Hernquist (2001) found that the fallback material with excess angular momentum circularises to form a disk on a typical viscous timescale:

$$t_\text{v} \sim 2.08\,\text{s} \times 10^3 \left(\frac{T_\text{c}}{10^6\,\text{K}} \right)^{-1} \left(\frac{r_\text{d}}{10^8\,\text{cm}} \right)^{1/2}, \qquad (3.7)$$

where T_c is the disk's central temperature and r_d is the circularisation radius of the disk. We assume the circularisation radius r_d to be equal to the Keplerian radius corresponding to the initial angular momentum of the fallback matter. Consider a mass element orbiting with a velocity v at the Keplerian radius r_d around a central neutron star of mass M_{NS}. The angular momentum density (i.e., the angular momentum per unit mass) associated with this mass element is $j = vr_\text{d}$. The Keplerian radius of the orbit is given by $r_\text{d} = GM_{\text{NS}}/v^2$. Therefore, by combining these two equations, we find that:

$$j = (GM_{\text{NS}}r_\text{d})^{1/2} \implies r_\text{d} = \frac{j^2}{GM_{\text{NS}}} \qquad (3.8)$$

Supernova simulations have shown that typical values for the angular momentum density of the fallback matter are around $j_{\text{fb}} \sim 10^{16-17}\,\text{cm}^2\,\text{s}^{-1}$ (Janka, Wongwathanarat, and Kramer, 2022), which corresponds to a circularisation radius of around $10^{6-8}\,\text{cm}$. In the following, we will always assume the fiducial values $T_\text{c} = 10^6\,\text{K}$ and $r_\text{d} = 10^8$ cm (see also Hameury et al., 1998).

The viscous timescale defined in Equation (3.7) determines the duration of an initial transient accretion phase characterised by a nearly constant accretion rate. Afterwards, as the supply of fallback matter decreases, the accretion rate into the disk itself declines as a power law. Furthermore, the disk starts to spread due to viscous effects. We follow these prescriptions and model the long-term time evolution of

the accretion rate and outer radius of the disk as (see Menou, Perna, and Hernquist, 2001; Ertan et al., 2009):

$$\dot{M}_{\mathrm{d}}(t) = \dot{M}_{\mathrm{d},0} \left(1 + \frac{t}{t_{\mathrm{v}}}\right)^{-\alpha}, \tag{3.9}$$

$$r_{\mathrm{out}}(t) = r_{\mathrm{d}} \left(1 + \frac{t}{t_{\mathrm{v}}}\right)^{\gamma}, \tag{3.10}$$

where the coefficients α and γ depend on the main mechanism determining the opacity of the disk. In particular, $\alpha = 19/16 \simeq 1.18$ and $\gamma = 3/8 \simeq 0.38$ if the disk opacity is dominated by electron scattering or $\alpha = 5/4 = 1.25$ and $\gamma = 1/2$ if the disk is dominated by Kramer's opacity (Cannizzo, Lee, and Goodman, 1990; Menou, Perna, and Hernquist, 2001; Ertan et al., 2009). In the following, we adopt intermediate values of $\alpha = 1.2$ and $\gamma = 0.44$.

In general, determining the fraction of fallback matter that eventually forms the disk is not trivial, since this depends on the progenitor's properties and on the supernova mechanism itself. Here, we consider a broad range of values for the disk initial accretion rate $\dot{M}_{\mathrm{d},0}$ between $10^{19-29}\,\mathrm{g\,s^{-1}}$. This is compatible with the disk accreting at a fraction of the overall supernova fallback rate as well as having a fraction of the total fallback mass M_{fb}.

From the accretion rate inside the disk we can compute an estimate of the accretion luminosity. Consider a mass element $\mathrm{d}M$ in the accretion disk around the neutron star. To fall from a Keplerian orbit at radius $r + \mathrm{d}r$ to an orbit of radius r, the mass element must lose gravitational potential energy. The virial theorem (Prialnik, 2000; Binney and Tremaine, 2008) dictates that half the potential energy is converted into additional kinetic energy. The remaining half is converted to heat due to viscous friction. The thermal energy released as the mass element moves inward will therefore be:

$$\mathrm{d}E_{\mathrm{th}} = \frac{1}{2}GM_{\mathrm{NS}}\mathrm{d}M\left(\frac{1}{r} - \frac{1}{r + dr}\right) \simeq \frac{1}{2}GM_{\mathrm{NS}}\mathrm{d}M\frac{\mathrm{d}r}{r^2}, \tag{3.11}$$

where we assume that $\mathrm{d}r \ll r$.

For a radiatively efficient disk we assume that the hot gas radiates its thermal energy as a black body at the same radius where the gravitational energy is liberated. The luminosity from an annulus in the disk characterised by the accretion rate $\dot{M}_{\mathrm{d}} = \mathrm{d}M/\mathrm{d}t$ is then given by:

$$\mathrm{d}L = \frac{\mathrm{d}E_{\mathrm{th}}}{\mathrm{d}t} = \frac{1}{2}GM_{\mathrm{NS}}\dot{M}_{\mathrm{d}}\frac{\mathrm{d}r}{r^2}. \tag{3.12}$$

The total luminosity of the accretion disk with inner and outer radii, r_{in} and r_{out} respectively, is found by integrating the luminosity over all annuli:

$$L_{\mathrm{acc}} = \int_{r_{\mathrm{in}}}^{r_{\mathrm{out}}} \frac{1}{2}\frac{GM_{\mathrm{NS}}\dot{M}_{\mathrm{d}}}{r^2}\mathrm{d}r = \frac{1}{2}GM_{\mathrm{NS}}\dot{M}_{\mathrm{d}}\left(\frac{1}{r_{\mathrm{in}}} - \frac{1}{r_{\mathrm{out}}}\right) \simeq \frac{GM_{\mathrm{NS}}\dot{M}_{\mathrm{d}}}{2r_{\mathrm{in}}}, \tag{3.13}$$

where we assumed that the mass accretion rate in the disk $\dot{M}_{\mathrm{d}}$ does not depend on the radius r and that $r_{\mathrm{in}} \ll r_{\mathrm{out}}$. If the flow of matter is constant throughout the disk up to the inner disk radius r_{in}, fiducial values $M_{\mathrm{NS}} = 1.4M_{\odot}$, $\dot{M}_{\mathrm{d}} = 10^{28}\,\mathrm{g\,s^{-1}}$ and $r_{\mathrm{in}} = 10^8$ cm result in luminosities up to $L_{\mathrm{acc}} \simeq 10^{46}\,\mathrm{erg\,s^{-1}}$. These luminosities by

far exceed the Eddington limit given by (Rybicki and Lightman, 1986):

$$L_{\text{Edd}} = \frac{4\pi G M_{\text{NS}} m_{\text{p}} c}{\sigma_{\text{T}}} \simeq 1.8 \times 10^{38} \, \text{erg s}^{-1} \left(\frac{M_{\text{NS}}}{1.4 \, M_\odot} \right), \quad (3.14)$$

where m_{p} is the proton mass and σ_{T} is the Thomson cross section and we naively assumed a disk composition of pure ionised hydrogen. It is generally expected that such super-critical inflows of matter produce winds and outflows that reduce the accretion rate in the inner disk region to values below the Eddington limit (Poutanen et al., 2007). We note however that in the presence of strong magnetic fields ($B \gtrsim 10^{14}$ G), the Eddington limit could be increased due to a reduction in the electron-scattering opacity (Canuto, Lodenquai, and Ruderman, 1971; Bachetti et al., 2014). In our model, we neglect this effect as it is relevant only in the very early stages of the evolution when the inner disk radius closely approaches the neutron star surface but it does not significantly affect the long-term dynamics outlined below. In a simplified scenario, we hence assume that the accretion rate at the inner disk radius r_{in} has to be limited by the Eddington accretion rate, given by $\dot{M}_{\text{Edd}} \simeq 2 L_{\text{Edd}} r_{\text{in}} / (G M_{\text{NS}}) \simeq 10^{18} (r_{\text{in}} / R_{\text{NS}}) \, \text{g s}^{-1}$. Therefore, we model the accretion rate at r_{in} as:

$$\dot{M}_{\text{d,in}}(t) = \begin{cases} \dot{M}_{\text{Edd}} & \text{if } \dot{M}_{\text{d}} \geq \dot{M}_{\text{Edd}}, \\ \dot{M}_{\text{d}}(t) & \text{if } \dot{M}_{\text{d}}(t) < \dot{M}_{\text{Edd}}. \end{cases} \quad (3.15)$$

In the following we discuss how the evolution of the disk and its influence on the spin-period evolution of the central neutron star depends on the complex interaction with the neutron star magnetosphere.

3.4 Period evolution in the presence of a fallback disk

Hereafter, we will study in detail the case where a disk forms successfully and interacts with the central neutron star. In the following for simplicity, we assume that the star's magnetic dipole moment is aligned with the rotation axis and that the magnetic field is crust-dominated and decays in time according to Equation (3.4). We also assume that the accretion disk's rotation is prograde with respect to the spin of the neutron star. We note that in order to observe pulsed radio emission from sources of this kind, a misalignment between the magnetic axis and the spin axis is, in principle, required. We do not study the impact of this effect in our calculations but point out that for small misalignment angles the evolutionary scenario should not differ substantially from the one studied in this work.

We consider two different cases depending on if the disk is able to influence the closed magnetosphere of the neutron star or not. If the accretion rate is sufficiently high, the in-falling matter is able to deform and penetrate the closed magnetosphere whose boundary can be roughly defined by the light cylinder radius r_{lc} (Equation (1.26)). In this case, we obtain the *magnetospheric radius* r_{m} which represents the distance from the star where the magnetic pressure equals the ram pressure of the accreted flow (Davidson and Ostriker, 1973; Elsner and Lamb, 1977; Ghosh and Lamb, 1979) through:

$$\frac{B^2}{8\pi} = \frac{1}{2} \rho v_{\text{ff}}, \quad (3.16)$$

where ρ is the mass density of the accreting matter and $v_{\text{ff}} = (2GM_{\text{NS}}/r)^{1/2}$ is its free-fall velocity at radius r. Assuming that the disk interacts with the magnetosphere at the magnetic equator, we can express the magnetic field in terms of the magnetic dipole moment as $B = \mu/r^3$ (see Equation (1.45)). Furthermore, by considering for simplicity spherical accretion we can express the accretion rate as $\dot{M}_{\text{d,in}} = 4\pi r^2 \rho v_{\text{ff}}$. By substituting this information into Equation (3.16) we can solve to find the magnetospheric radius:

$$r_{\text{m}} = \xi \left(\frac{\mu^4}{2GM_{\text{NS}}\dot{M}_{\text{d,in}}^2} \right)^{1/7}, \qquad (3.17)$$

where $\xi \sim 0.5$ is a corrective factor that takes into account that the accretion disk has a non-spherical geometry (Long, Romanova, and Lovelace, 2005; Bessolaz et al., 2008; Zanni and Ferreira, 2013). To relate the stellar magnetic moment to the time-dependent magnetic field at the pole $B_{\text{p}}(t)$ (assuming a dipolar structure for the magnetosphere), we can use Equation (1.44) to obtain $\mu = B_{\text{p}}(t)R_{\text{NS}}^3/2$. Note that the magnetospheric radius evolves in time due to the decay of both the disk accretion rate and the magnetic field. Inside the magnetospheric radius, the plasma dynamics are dominated by the magnetic field such that the accretion flow is forced to corotate with the star's closed magnetosphere. Under such conditions, the magnetospheric radius roughly determines the inner edge of the accretion disk so that we can assume $r_{\text{in}} \sim r_{\text{m}}$.

If the accretion rate is sufficiently low, the magnetic pressure of the closed magnetosphere is able to keep the accretion flow outside the light cylinder. However beyond r_{lc}, the dipole configuration breaks down and the magnetic field lines open up. As a consequence, they are no longer able to exert significant pressure on the accreted plasma. Under these conditions, we expect the inner radius r_{in} of the disk to be roughly equal to r_{lc}. Thus, in general we adopt the same prescription as Yan, Perna, and Soria (2012) and assume $r_{\text{in}} \simeq \min(r_{\text{m}}, r_{\text{lc}})$.

Note that a condition for the disk to form and remain active is that the disk's outer radius satisfies $r_{\text{out}} > r_{\text{in}}$. Otherwise, the stellar magnetic field would completely disrupt the disk, restoring a configuration where the neutron star's spin-down is determined by dipolar losses only.

Another critical lengthscale is the *corotation radius*

$$r_{\text{cor}} = \left(\frac{GM_{\text{NS}}}{\omega^2} \right)^{1/3}, \qquad (3.18)$$

which represents the distance at which the gravitational pull of the neutron star balances the centrifugal force for a test mass that is corotating with the star at the spin frequency ω.

The position of the magnetospheric radius with respect to the light cylinder radius and the corotation radius determines the total torque N_{tot} exerted on the star, which in turn drives the time evolution of the spin period according to:

$$I_{\text{NS}}\dot{\omega} = N_{\text{tot}}. \qquad (3.19)$$

Following Piro and Ott (2011) and Metzger, Beniamini, and Giannios (2018) the total torque can be modelled by the following equation:

$$
\begin{aligned}
N_{\text{tot}} &= N_{\text{acc}} + N_{\text{dip}} \\
&= \dot{M}_{\text{d,in}} r_{\text{in}}^2 \left[\Omega_K(r_{\text{in}}) - \omega \right] - I_{\text{NS}} \left(\frac{r_{\text{lc}}}{r_{\text{in}}} \right)^2 \beta B^2 \omega^3 \\
&= \begin{cases} \dot{M}_{\text{d,in}} r_{\text{lc}}^2 \left[\Omega_K(r_{\text{lc}}) - \omega \right] - I_{\text{NS}} \beta B^2 \omega^3 & \text{if } r_{\text{m}} > r_{\text{lc}}, \\ \dot{M}_{\text{d,in}} r_{\text{m}}^2 \left[\Omega_K(r_{\text{m}}) - \omega \right] - I_{\text{NS}} \left(\frac{r_{\text{lc}}}{r_{\text{m}}} \right)^2 \beta B^2 \omega^3 & \text{if } r_{\text{m}} \le r_{\text{lc}}. \end{cases}
\end{aligned}
\tag{3.20}
$$

where N_{dip} accounts for the electromagnetic torque of the magnetosphere while N_{acc} accounts for the torque exerted by the accretion process and $\Omega_K(r) = \left(GM_{\text{NS}}/r^3 \right)^{1/2}$ is the Keplerian orbital angular velocity at radius r.

Figure 3.3 shows an example of a solution of Equation (3.19) for a pulsar interacting with a fallback disk. We consider an initial spin period $P_0 = 10\,\text{ms}$, initial magnetic field $B_0 = 4 \times 10^{14}\,\text{G}$ and an initial disk accretion rate $\dot{M}_{\text{d,0}} = 10^{24}\,\text{g}\,\text{s}^{-1}$. Note that for this choice of initial parameters the disk can form and stay active since $r_{\text{out}} > r_{\text{in}}$ at all times. Furthermore, in the early phases, accretion at the inner radius of the disk is limited by the Eddington limit (see Section 3.3).

Different regimes are present depending on the relative ordering of the three radii defined above. In particular, with reference to Figures 3.2 and 3.3, we distinguish the three following phases.

Ejector phase: for $r_{\text{m}} > r_{\text{lc}} > r_{\text{cor}}$, the accreted material remains at the boundary of the closed magnetosphere and does not influence the neutron star's internal dynamics, i.e., the star spins down mainly due to dipolar electromagnetic torques, that is $N_{\text{tot}} \sim N_{\text{dip}}$ (top panel in Figure 3.2). After a phase of duration t_{em} where the spin period stays constant, P starts to increase $\propto t^{1/2}$ (see Equation (3.2)). As a consequence, the characteristic radii r_{lc} and r_{cor} increase as $t^{1/2}$ and $t^{2/3}$, respectively. This phase is commonly referred to as the ejector phase (shaded yellow region in Figure 3.3). The mechanism responsible for radio emission can be active and the neutron star could be observed as a radio pulsar. Morover, as $\dot{M}_{\text{d}}$ decreases and eventually becomes sub-Eddington, the magnetospheric radius grows slower than the other two critical radii ($\propto t^{2\alpha/7} \simeq t^{0.34}$ for $\alpha = 1.2$). Thus, r_{m} will eventually cross the light cylinder.

Propeller phase: for $r_{\text{lc}} > r_{\text{m}} > r_{\text{cor}}$, the accretion flow is able to penetrate inside the closed magnetosphere (middle panel in Figure 3.2). As it reaches the magnetospheric radius, the plasma flow is forced to corotate with the magnetosphere at super-Keplerian speeds causing it to be ejected due to centrifugal forces. This introduces a viscous torque that spins down the star very efficiently, a phase commonly referred to as the propeller (shaded orange region in Figure 3.3). In this case, we have $\omega > \Omega_K(r_{\text{m}})$. The electromagnetic torque N_{dip}, i.e., the second term in Equation (3.20), is also influenced by accretion. When the accretion flow penetrates inside the closed magnetosphere, the fraction of the magnetic field lines, which connect the star to the disk, is forced to open and as a consequence, the polar cap region containing open field lines expands. This enhances the spin-down torque caused by the magnetosphere (Parfrey, Spitkovsky, and Beloborodov, 2016; Metzger, Beniamini, and Giannios, 2018). In particular Parfrey, Spitkovsky, and Beloborodov (2016) argue that for $r_{\text{m}} < r_{\text{lc}}$, the magnetic flux through the expanded polar cap increases
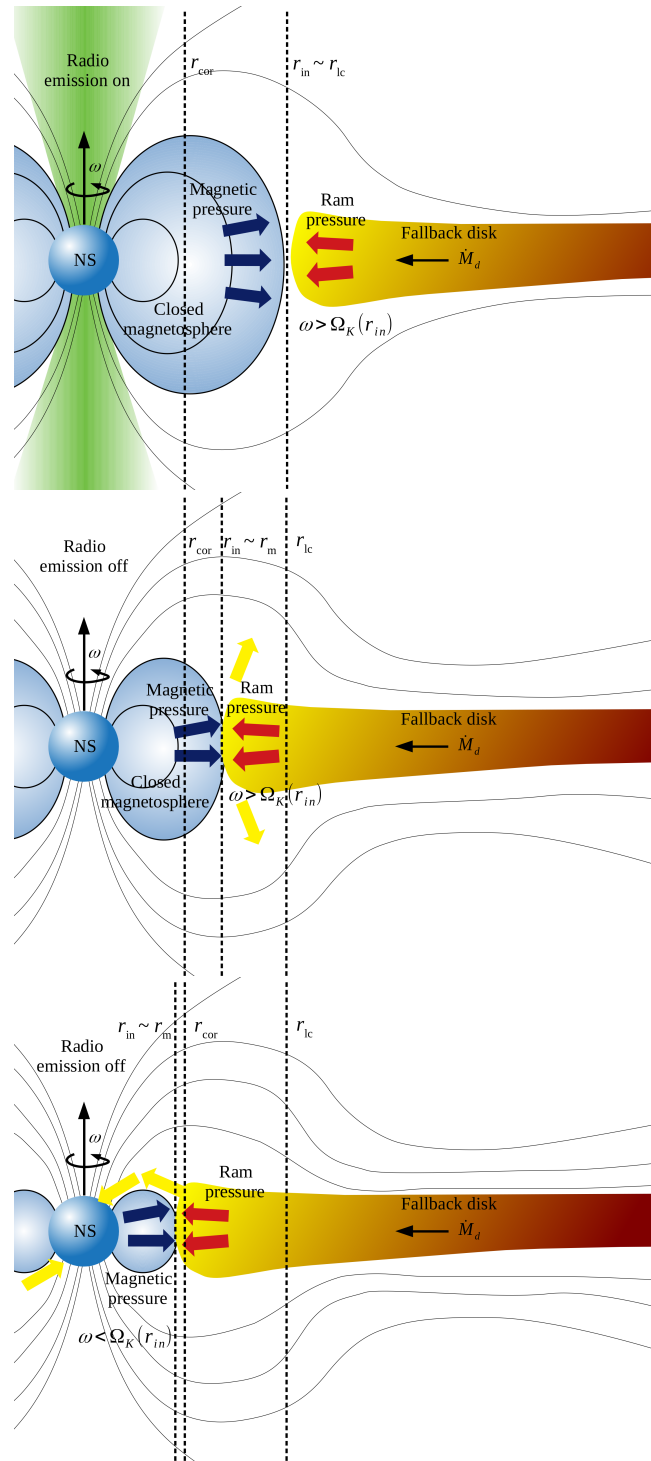
FIGURE 3.2: Sketch showing the three phases a neutron star with a fallback disk system can experience during its evolution. During the ejector phase (*top panel*) the fallback disk is confined outside the light cylinder and the neutron star spin down mainly by dipolar electromagnetic losses. In the propeller phase (*middle panel*) the fallback disk penetrate inside the closed magnetosphere but at the r_{in} it finds a centrifugal barrier. The neutron star spin down mainly by propeller torque. In the direct accretion phase (*bottom panel*) the fallback disk penetrates inside the closed magnetosphere and it is super-Keplerian at r_{in}. The neutron star spins up and accretes on the surface.
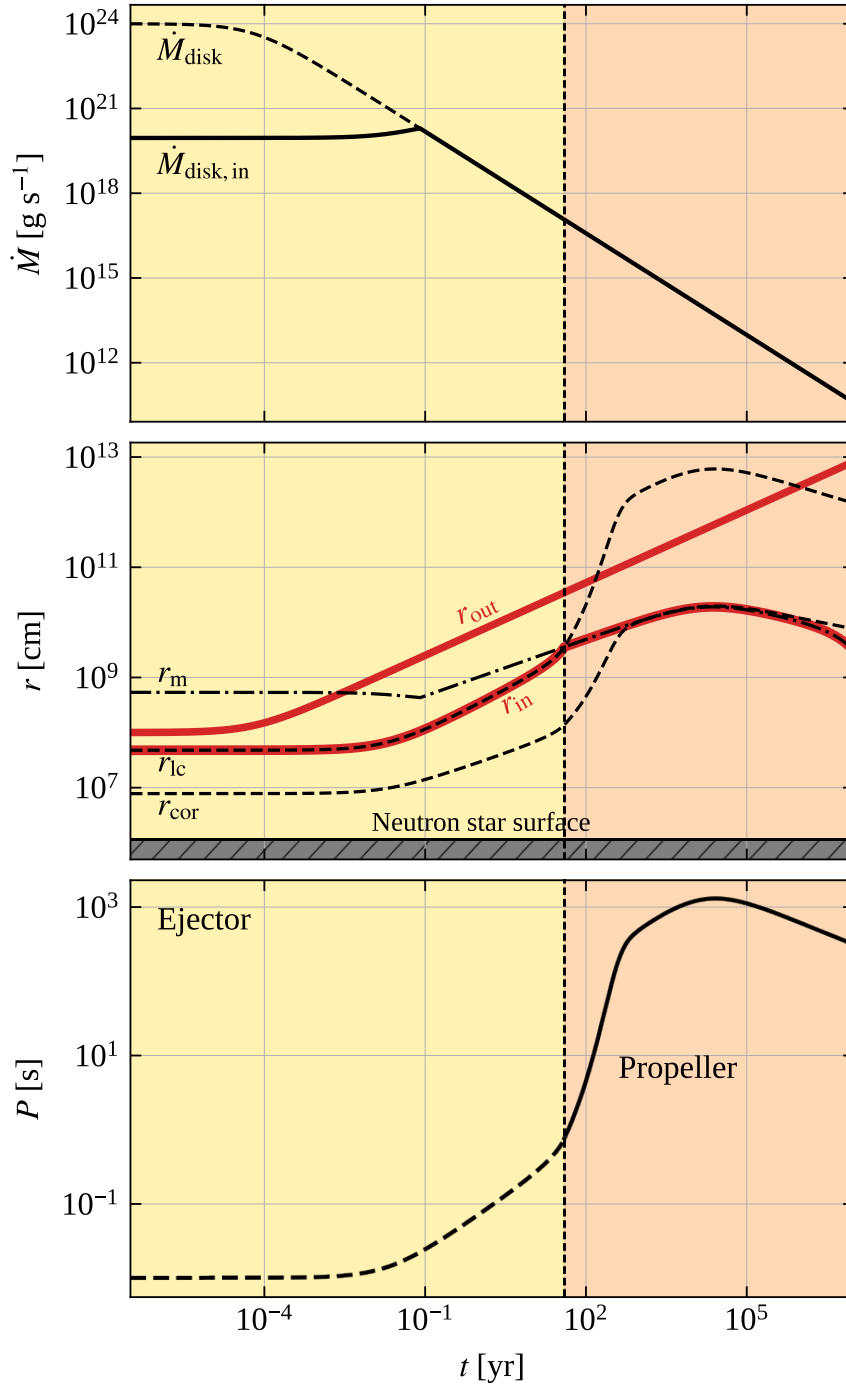
FIGURE 3.3: Example of neutron star spin-down in the presence of a fallback disk on a timescale of 10^7 yr. The top panel shows the total disk accretion $\dot{M}_d$ and the Eddington-limited accretion rate at the inner radius $\dot{M}_{d,in}$. The middle panel illustrates the evolution of the three critical radii r_{cor}, r_m, r_{lc}. The neutron star radius is also indicated as a reference. The evolution of the disk's inner and outer radii is highlighted by the red lines. The bottom panel shows the resulting time evolution of the spin period. We assume an initial spin period $P_0 = 10$ ms, initial magnetic field $B_0 = 4 \times 10^{14}$ G and an initial disk accretion rate $\dot{M}_{d,0} = 10^{24}$ g s^{-1}. We also highlight the duration of the ejector and propeller phases by shading the background in yellow and orange, respectively.

by a factor $\sim (r_{\mathrm{lc}}/r_{\mathrm{m}})$. As a consequence, since the dipole spin-down torque is proportional to the square of the magnetic flux through the open field-line region, N_{dip} is enhanced by a factor $\sim (r_{\mathrm{lc}}/r_{\mathrm{m}})^2$. Moreover, since the radio emission is associated with magnetospheric currents and pair production (Beloborodov, 2008; Philippov, Timokhin, and Spitkovsky, 2020, see also Section 1.7 and Section 1.9.3), the mechanism generating the radio emission is likely perturbed or even stopped as the closed-magnetosphere geometry is disturbed by the accretion flow. This effectively switches off the radio-loud nature of these sources (Li, 2006). Therefore, we expect a neutron star in the propeller regime unlikely to be observable as a radio pulsar. During this propeller phase, the spin frequency decreases over time and eventually $\omega \sim \Omega_K(r_{\mathrm{m}})$ (i.e., when $r_{\mathrm{cor}} \sim r_{\mathrm{m}}$), so that the net torque exerted on the neutron star vanishes and a spin equilibrium is reached. From this point onward, further evolution of the spin period will be regulated mainly by time variation of the accretion rate in the inner disk $\dot{M}_{\mathrm{d,in}}$ and the stellar magnetic field B. For example, the decay of the magnetic field strength causes the slow decrease of the magnetospheric radius at late times. This drives a progressive shift of the spin-equilibrium radius towards the neutron star surface, causing the neutron star to slightly spin up; in Figure 3.3, this happens when the magnetic field starts to decay after a Hall timescale, i.e., after $\sim 10^4$ yr.

Direct accretion phase: for $r_{\mathrm{lc}} > r_{\mathrm{cor}} > r_{\mathrm{m}}$, we have $\omega < \Omega_K(r_{\mathrm{m}})$ (bottom panel in Figure 3.2). The accretion flow still manages to penetrate inside the closed magnetosphere but as it reaches the magnetospheric radius, the plasma is forced to corotate with the magnetosphere at sub-Keplerian speeds. In other words, at the boundary defined by the magnetospheric radius, the accreted plasma is orbiting faster than the neutron star magnetosphere. This introduces a viscous torque that transfers angular momentum from the plasma flow to the star and tends to spin up the star. Besides this, the enhanced electromagnetic spin-down torque described above is still acting, opposing the spin-up due to accretion. Therefore, the evolution in this regime is controlled by the relative strength of N_{acc} and N_{dip}. In the example in Figure 3.3, this phase is experienced at very late times, i.e. $\sim 10^6$ yr, when the magnetic field has decayed so much that the star exits the spin equilibrium and the magnetospheric radius becomes smaller than the corotation radius.

From these considerations, we observe that the parameters that mainly regulate the evolution of a pulsar surrounded by a fallback disk are the initial stellar magnetic field B_0 and the initial disk accretion rate $\dot{M}_{\mathrm{d,0}}$. The latter determines the time at which the accretion rate affecting the compact object becomes sub-Eddington and starts decreasing. In contrast, the value of the neutron star's initial spin period only affects the early evolution stages, while the long-term evolution of the neutron star is almost insensitive to the value of P_0.

Figures 3.4 and 3.5 show several examples of evolutionary curves of the spin period. We always assume an initial spin period $P_0 = 10$ ms. In particular, Figure 3.4 shows the spin-period evolution for several values of the initial magnetic field B_0 and varying initial disk accretion rate $\dot{M}_{\mathrm{d,0}}$, while Figure 3.5 shows the evolution curves for several values of initial disk accretion rate and varying initial magnetic field. In the early phases, the accretion at the inner radius of the disk is limited by the Eddington limit. For each simulated evolutionary curve we check that the disk's outer radius is always greater than the inner radius guaranteeing that the disk can form and influence the rotational evolution of the neutron star. In general,
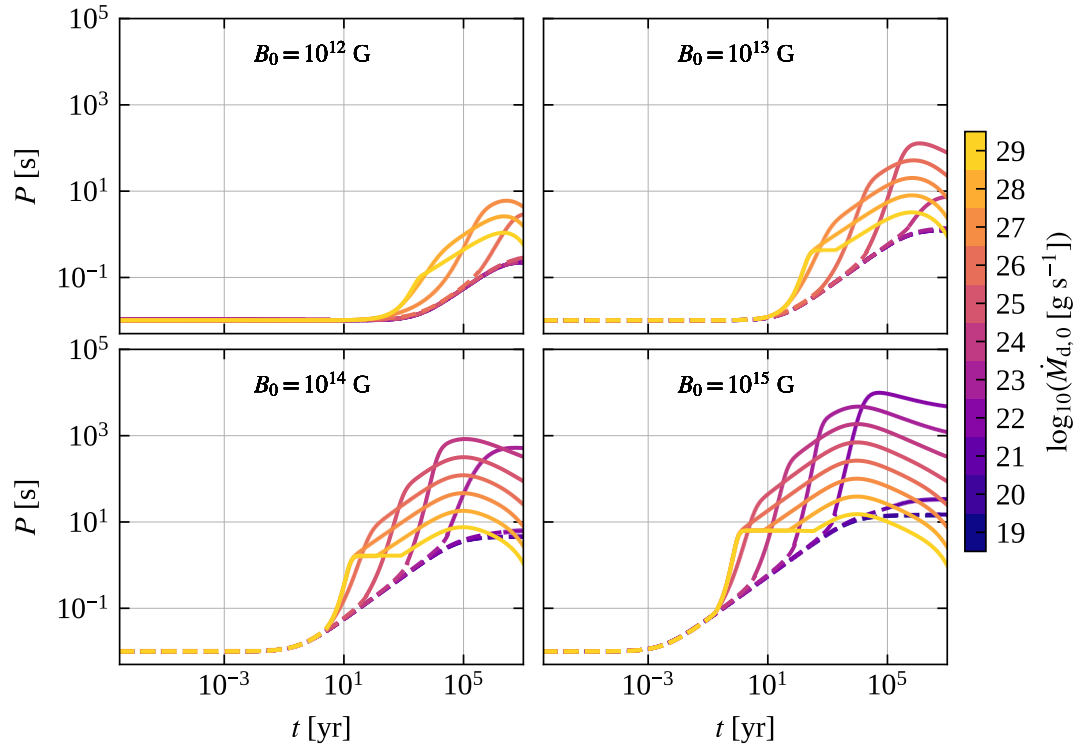
FIGURE 3.4: Example curves showing the time evolution of the spin period for a pulsar interacting with a fallback disk on a timescale of 10^7 yr, for different assumptions on the initial B_0 field strength, and varying disk fallback rate $\dot{M}_{d,0}$. The dashed portion of the curves indicates when the neutron star is in the radio-loud ejector phase, i.e., when $r_m > r_{lc}$, while the solid portion indicates when the neutron star is in the radio-quiet propeller stage, i.e., when $r_m < r_{lc}$.
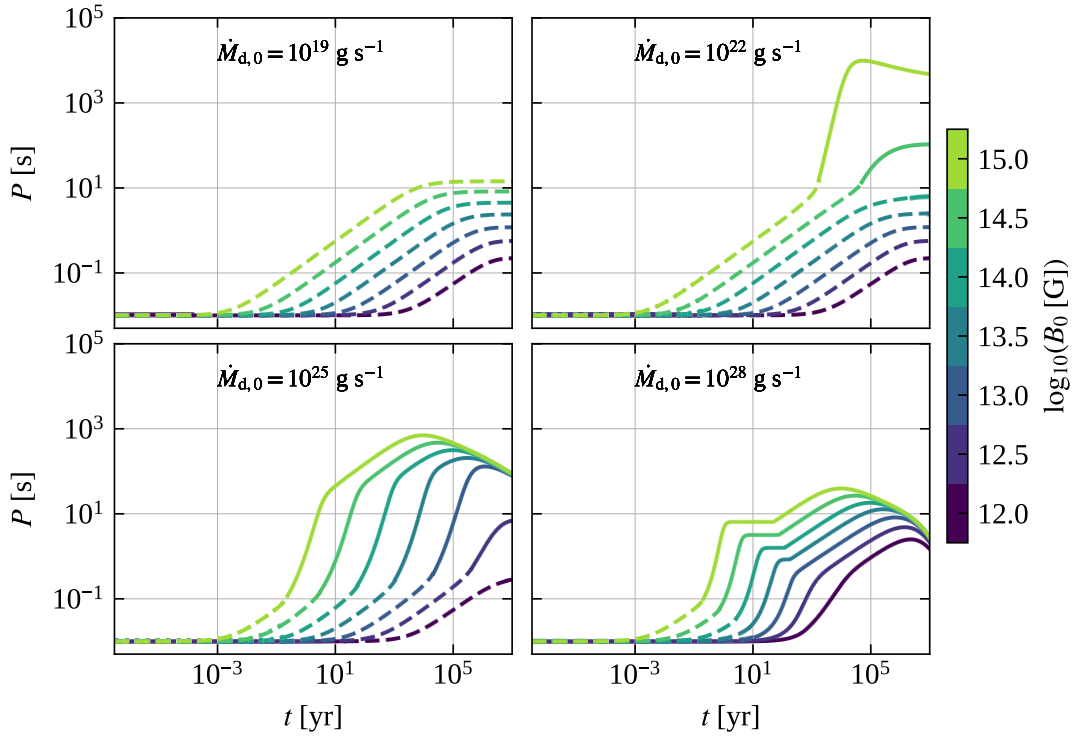
FIGURE 3.5: Example curves showing the time evolution of the spin period for a pulsar interacting with a fallback disk on a timescale of 10^7 yr, for different assumptions of the disk fallback rate $\dot{M}_{d,0}$, and varying the initial magnetic field B_0. The dashed portion of the curves indicates when the neutron star is in the radio-loud ejector phase, i.e., when $r_m > r_{lc}$, while the solid portion indicates when the neutron star is in the radio-quiet propeller stage, i.e., when $r_m < r_{lc}$.

it can be noted that for low values of the initial disk accretion rate ($< 10^{22}\,\mathrm{g\,s^{-1}}$), the neutron star remains in the ejector phase for its entire evolution, independent of the value of its initial magnetic field. In contrast, higher accretion rates and higher magnetic fields allow the star to enter the propeller phase at earlier times. However, the propeller regime is most effective at spinning down the neutron star to periods $\gtrsim 10\,\mathrm{s}$ only for relatively strong magnetic fields $\gtrsim 10^{13}\,\mathrm{G}$ and intermediate initial disk accretion rates in the range $10^{22-27}\,\mathrm{g\,s^{-1}}$.

We note that a neutron star that enters the propeller phase (and subsequently reaches spin equilibrium) will remain in this state until an abrupt change in the disk accretion rate occurs. For example, if the accretion rate in the disk suddenly drops, a neutron star can exit the propeller phase and enter the ejector phase again. In this case, the neutron star can transition from a faint X-ray source (due to thermal emission from material accreted onto the magnetosphere) to a standard rotation-powered radio pulsar or potentially radio-loud magnetar (see Section 3.9 for a detailed explanation of how this transition can occur and a description of the expected X-ray and radio luminosity in the different fallback accretion states).

3.5 Accretion from spherical fallback

If the fallback matter does not possess sufficient angular momentum to form a disk, the fallback will proceed quasi-spherically. Even in this case, the neutron star could experience different accretion phases depending on the magnitude of the fallback rate. If the matter inflow is radiatively inefficient, accretion could proceed at super-Eddington rates, especially in the early stages. For such high fallback rates, the accretion flow is likely able to penetrate and squeeze the proto-neutron star magnetosphere, causing an initial phase of direct accretion onto the surface. This might also result in the burial of the magnetic field (see for example Taam and van den Heuvel, 1986; Li et al., 2021; Lin et al., 2021), a scenario that (combined with the subsequent secular re-emergence of the magnetic field) has been invoked to explain the observed properties of Central Compact Objects (CCOs) (Halpern and Gotthelf, 2010; Fu and Li, 2013; Ho, 2015; Zhong et al., 2021); a class of young, generally weak-field neutron stars found close to the centres of supernova remnants (introduced in Section 1.13). After this initial direct accretion phase, as the fallback rate decreases in time, the neutron star could, in principle, enter a propeller phase that causes the star to spin down as in the disk scenario. However, in the case of quasi-spherical accretion, the fallback episode is expected to last at most a free-fall timescale $t_{\mathrm{ff}} \sim (G\rho)^{-1/2}$, determined by the density ρ of the infalling outer layers of the progenitor participating in the fallback. In general, t_{ff} could reach at most $\sim 10^7\,\mathrm{s}$ for red supergiant progenitor stars, whose envelopes have typical densities $\rho \sim 10^{-7}\,\mathrm{g\,cm^{-3}}$ (Metzger, Beniamini, and Giannios, 2018). Therefore, it is unlikely that a propeller phase acting on such short timescales (of the order of $\sim 1\,\mathrm{yr}$) could result in equally long spin periods as the disk scenario.

3.6 The 18-min periodic radio transient GLEAM-X J162759.5-523504.3

During the GaLactic and Extragalactic All-sky MWA eXtended survey (GLEAM-X) (Hurley-Walker et al., 2022b) with the Murchison Widefield Array (MWA), a peculiar periodic radio transient has been discovered displaying a periodicity of 1091 s
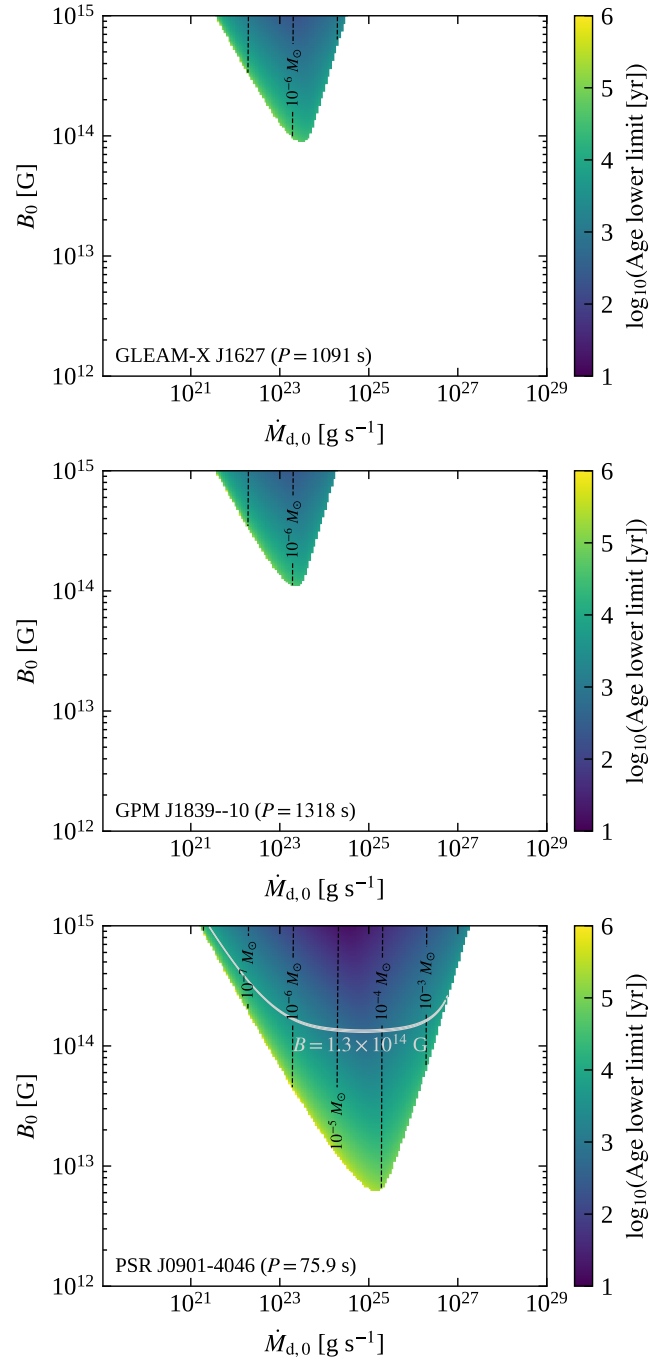
FIGURE 3.6: Long-period pulsars in the supernova fallback-accretion picture. The coloured regions show the values of the initial magnetic field B_0 and disk fallback rate $\dot{M}_{d,0}$ that allow the neutron star to reach a spin period of 1091 s for GLEAM-X J1627 (top left panel), 1318 s for GPM J1839–10 (top right panel) and 75.9 s for PSR J0901-4046 (bottom panel) in less than 10^7 yr. The colour code indicates the time at which the neutron stars have reached their respective periods and can be interpreted as a lower limit on the source's current age. The contour lines indicate the total fallback mass that has been accreted by the disk in the same time interval. For PSR J0901-4046, the gray contour line represents where the magnetic field value at $t =$ Age is equal to the estimated magnetic field of PSR J0901-4046 from the spin-down formula.

(Hurley-Walker et al., 2022a). The source was detected during two radio outburst periods in January and March 2018, displaying a very variable flux (going from undetected to values as high as $20 - 50\,\text{Jy}$), $\sim 5\%$ duty cycle, 90% linear polarisation, and a very spiky and variable pulse profile.

From a detailed timing analysis, a dispersion measure of DM $= 57 \pm 1\,\text{pc}\,\text{cm}^{-3}$ was calculated, converting to a distance of $1.3\,\text{kpc}$ according to the Galactic electron-density model of Yao, Manchester, and Wang (2017). The period derivative was loosely constrained to $\dot{P} < 1.5 \times 10^{-9}\,\text{s}\,\text{s}^{-1}$. Assuming that the source is an isolated neutron star spinning down due to the classical electromagnetic dipole formula, this gives a relatively weak constraint on the dipolar magnetic field of $B_{\text{dip}} < 10^{17}\,\text{G}$ through Equation (1.46). We can also obtain an upper limit on the spin-down power $\dot{E}_{\text{rot}} = I_{\text{NS}}\omega\dot{\omega} < 10^{28}\,\text{erg}\,\text{s}^{-1}$. If we roughly estimate an isotropic radio luminosity from the flux detected during the outburst phase and the distance estimated from the DM, we obtain a value of $\sim 10^{31}\,\text{erg}\,\text{s}^{-1}$. From this, we deduce that the spin-down alone is insufficient to power these very bright radio flares. The radio characteristics, such as the large flux variability, the radio outburst activity, and the high linear polarisation seem analogous to those observed in radio-active magnetars, whose activity is believed to be triggered and powered by the evolution and decay of their strong magnetic fields (see also Section 1.13 and Thompson and Duncan, 1995; Duncan and Thompson, 1996). However, GLEAM-X J1627's exceptionally long spin period would make this source stand out among them.

As shown in Figure 3.1, assuming GLEAM-X J1627 is indeed a magnetar that has spun down via dipolar losses alone, in the limiting case of a constant magnetic field, the magnetar would require an age $\gtrsim 10^6(10^8)\,\text{yr}$ and a very strong magnetic field of $10^{16}(10^{15})\,\text{G}$ to reach its spin period of $1091\,\text{s}$. As argued in Section 3.2, sustaining such a strong magnetic field over this lifetime is difficult to reconcile with crustal magnetic-field evolution models that predict field decay on the Hall timescale; in this case $\sim 10^3\,\text{yr}$. For comparison, SGR 0418+5729 is most likely the oldest magnetar detected so far with a characteristic age of $\sim 10^7\,\text{yr}$ and has an inferred magnetic field of $\sim 10^{13}\,\text{G}$ (Magnetar Outburst Online Catalogue http://magnetars.ice.csic.es/ Coti Zelati et al., 2018). Moreover, if the star's magnetic energy reservoir has decayed in time, it becomes challenging to explain the current magnetar-like activity observed from this source.

A more promising explanation for GLEAM-X J1627 that requires less extreme conditions is a magnetar that has experienced accretion from a fallback disk soon after the supernova. As outlined in §3.3, a magnetar surrounded by a fallback disk will pass through the propeller phase and spin down very efficiently on short timescales. To study this scenario for GLEAM-X J1627, we first fix the initial spin period to $P_0 = 10\,\text{ms}$ (remember that as long as $P \gg P_0$, P_0 has very little influence on the long-term evolution); for the neutron star mass and radius we adopt the fiducial values $M_{\text{NS}} = 1.4\,M_\odot$ and $R_{\text{NS}} = 11\,\text{km}$. By varying the two parameters B_0 and $\dot{M}_{\text{d,0}}$ and using Equation (3.20) to determine the torque acting on the star in the different stages, we can numerically integrate Equation (3.19) in time. This allows us to find those parameter combinations that lead to a spin-down evolution reaching a period of at least $1091\,\text{s}$. In what follows, we consider a maximum time of $10^7\,\text{yr}$ for the evolution. The motivation for this limit is two-fold. Firstly, it ensures that the spin-period evolution curves reach their maximum values during the propeller phase before magnetic field decay enters into play (see the discussion of Figures 3.3–3.5 in Section 3.4). Secondly, a limit of $10^7\,\text{yr}$ ensures that we encompass all age estimates of currently known magnetars (see the Magnetar Outburst Online Catalogue http://magnetars.ice.csic.es/ Coti Zelati et al. (2018)).

The top panel in Figure 3.6 shows the parameter space of B_0 and $\dot{M}_{d,0}$, where we choose an initial magnetic field between 10^{12-15} G and an initial disk accretion rate between 10^{19-29} g s^{-1}. The coloured region represents the combination of parameters that guarantees the neutron star to reach a spin period of 1091 s. The colour code indicates the age at which the neutron star reaches the desired period for the first time. The vertical contour lines show the total fallback mass that has been accreted into the disk over the same time interval (calculated by integrating Equation (3.9) in time). In this scenario, we observe that (thanks to the propeller phase) a magnetar with a magnetic field of around 10^{14} G can reach a spin period of 1091 s on a relatively short time scale of about 10^{3-5} yr for an initial disk accretion rate of around 10^{23} g s^{-1} and a total accreted mass of $\sim 10^{-6} M_\odot$. Note that since GLEAM-X J1627 has been observed to emit periodic radio signals, disk accretion must have ceased to allow for the radio activity to be restored (see §3.9). Assuming this is the case, the ages shown in Figure 3.6 should be interpreted as the times at which the neutron star exited the propeller phase and shifted to standard dipolar spin-down; i.e., they represent lower limits on the real age of the observed source.

3.7 The 21-min periodic radio transient GPM J1839-10

During recent monitoring of the Galactic plane using the MWA and subsequent follow-up observations with the Australia Telescope Compact Array (ATCA), Parkes Murriyang radio telescope, the Australian Square Kilometre Array Pathfinder (ASKAP) and MeerKAT another mysterious periodic radio source with a periodicity of 1318 s has been discovered (Hurley-Walker et al., 2023). Further searches in archival data from the Very Large Array (VLA), the Low Band Ionospheric and Transient Experiment (VLITE) and the Giant Metrewave Radio Telescope (GMRT) demonstrated that this source has been active for at least 33 years. As for GLEAM-X J1627 the morphology and brightness of the observed pulses is very variable, going from undetected to maximum flux densities in the range of 0.1 Jy $- 10$ Jy with linear polarisation varying between 10% and 100% and a duty cycles as large as 20%. The timing analysis allowed an accurate estimation of the dispersion measure DM $= 273.5 \pm 2.5$ pc cm^{-3} which corresponds to a distance of 5.7 ± 2.9 kpc according to the Galactic electron density model of Yao, Manchester, and Wang (2017). The long-duration activity of this source allowed us to infer a stringent constraint on its period derivative $\dot{P} \lesssim 3.6 \times 10^{-13}$ s s^{-1}. Assuming the source to be an isolated neutron star spinning down via dipole radiation, we find upper limits on the polar magnetic field $B_{\rm dip} \lesssim 1.2 \times 10^{15}$ G and on the rotational power $\dot{E}_{\rm rot} \lesssim 8.4 \times 10^{24}$ erg s^{-1}. A rough estimate of the radio luminosity obtained from the observed flux densities and the estimated distance leads to $L \sim 10^{28}$ erg s^{-1}. As for GLEAM-X J1627 source, the estimated rotational energy loss is not enough to power the observed radio luminosity. Furthermore assuming a magnetar is producing the emission of these bright pulses, extreme values of the magnetic field are necessary to reach such long spin periods (see Figure 3.1).

We extend the analysis of Ronchi et al. (2022) and we also study the period evolution in the fallback scenario for GPM J1839–10. We solve the spin evolution equation Equation (3.19) for different combinations of the two parameters B_0 and $\dot{M}_{d,0}$ and determine the parameter space that allows GPM J1839–10 to reach a spin period of at least $P = 1318$ s within $< 10^7$ yr. We show the results in the middle panel of Figure 3.6. We notice that the allowed region in the $B_0 - \dot{M}_{d,0}$ plane is very similar

to that for GLEAM-X J1627. We again highlight that GPM J1839–10's disk accretion must have stopped in order to allow the re-activation of the radio-pulsar mechanism.

3.8 The 76 seconds magnetar: PSR J0901-4046

The long-period pulsar PSR J0901-4046 was recently discovered by the South African radio telescope MeerKAT. The pulsar manifests a periodicity of $P = 75.9\,\mathrm{s}$ with a period derivative of $\dot{P} = 2.25 \times 10^{-13}\,\mathrm{s\,s^{-1}}$ (Caleb et al., 2022). From the timing analysis, a dispersion measure of $\mathrm{DM} = 52 \pm 1\,\mathrm{pc\,cm^{-3}}$ has been derived which translates into a distance of $\sim 0.4\,\mathrm{kpc}$. Assuming that this neutron star is spinning down simply by dipolar losses, the strength of the dipolar magnetic field can be computed to $B_{\mathrm{dip}} \simeq 1.3 \times 10^{14}\,\mathrm{G}$ and its spin-down power to $\dot{E}_{\mathrm{rot}} \simeq 2 \times 10^{28}\,\mathrm{erg\,s^{-1}}$. Such a strong field is in the typical range of observed magnetars (see Figure 1.12).

Using similar considerations as for GLEAM-X J1627 and GPM J1839–10 and assuming that PSR J0901-4046 is an isolated neutron star, which has spun down due to electromagnetic dipolar losses alone, we can reproduce the observed period, if its magnetic field remained almost constant over a lifetime of almost $10^7\,\mathrm{yr}$ (the source's characteristic age is $\tau_c = P/(2\dot{P}) \simeq 5.4\,\mathrm{Myr}$). If we instead consider the fallback disk scenario, we can relax these conditions. As before, we solve the spin evolution equation Equation (3.19) for different combinations of the two parameters B_0 and $\dot{M}_{\mathrm{d},0}$ and determine the parameter space that allows PSR J0901-4046 to reach a spin period of at least $P = 75.9\,\mathrm{s}$ within $< 10^7\,\mathrm{yr}$. The results are reported in the bottom panel of Figure 3.6. Also for PSR J0901-4046, disk accretion must have stopped in order for the radio emission mechanism to be restored. From Figure 3.6, we deduce that observational characteristics can be explained within our model if this happened at an age of around $10^{3-4}\,\mathrm{yr}$, which we again interpret as a lower limit on the pulsar's true age. This constraint on the age is in line with other radio-emitting magnetars or high-B radio pulsars that show similar radio emission. Furthermore in the case of PSR J0901-4046, an initial disk accretion rate of around $10^{24}\,\mathrm{g\,s^{-1}}$, and a total accreted mass of $\sim 10^{-5}M_\odot$ are able to explain the observed period.

3.9 Discussion

3.9.1 Re-establishing the radio emission after the propeller

One key piece that we have to discuss is how a neutron star, that enters the propeller phase to experience effective spin-down, can exit this phase again in order to be observed as a radio-loud object. This is a crucial requirement, given that we need to explain the radio detection of GLEAM-X J1627, GPM J1839–10 and PSR J0901-4046. As briefly mentioned in §3.3, the transition from the propeller to the ejector phase can be naturally explained by an abrupt drop in the accretion rate. As a result, the magnetospheric radius would move beyond the light cylinder, providing the condition for the radio emission to be reactivated.

Several factors could play a role in causing such a drop in the accretion flow. For example, since magnetars are very active young neutron stars that undergo outburst and flaring episodes (Coti Zelati et al., 2018; Esposito, Rea, and Israel, 2021), the energy released in such events could cause the disk to unbind or be completely disrupted. This would subsequently stop accretion and radio emission could be reinstated.

Another explanation could be simply that the fallback disk runs out of matter. During the propeller phase, the material that reaches the magnetospheric radius is ejected due to the centrifugal barrier. If the ejected material possesses a velocity superior to the escape velocity, it will become unbound from the system; otherwise, it will fall back and be reprocessed inside the accretion disk. As the matter feeding the disk from the supernova fallback is not replenished, we expect the accretion disk to eventually be completely consumed if the propeller is efficient at unbinding matter from the system (see for example Ekşi, Hernquist, and Narayan, 2005; Romanova et al., 2005).

Another possibility is that the fallback disk, which itself evolves with time, undergoes a thermal ionisation instability, which has been outlined in detail in Mineshige, Nomoto, and Shigeyama (1993), Menou, Perna, and Hernquist (2001), Ertan et al. (2009), and Liu and Li (2015). As the disk spreads and the accretion rate decays with time, energy dissipation decreases and the disk gradually cools down. In particular, Menou, Perna, and Hernquist (2001) argued that as the disk accretion rate falls below a critical value of around $10^{15} \, \mathrm{g \, s^{-1}}$ and the temperature in the outermost part of the disk drops below $\sim 10^4 \, \mathrm{K}$, the recombination of free electrons with heavy nuclei in the plasma is triggered. This transition alters the corresponding magnetic and viscous properties of the disk, reducing the efficiency of angular momentum transfer and eventually stopping accretion. As the disk evolves further in time, the recombination front propagates from the outermost regions inwards so that the active region of the disk shrinks. Eventually, accretion onto the neutron star will halt completely if the disk becomes totally neutralised. Menou, Perna, and Hernquist (2001) have argued that this transition occurs at around $10^{3-4} \, \mathrm{yr}$ when the outer disk radius is $\sim 10^{10-11} \, \mathrm{cm}$. This timescale is comparable to the ones we require in our fallback accretion scenario for neutron stars to reach spin periods $> 10 \, \mathrm{s}$ (see Figure 3.6). However, it has also been suggested that the irradiation from the central source may prevent the disk from becoming completely neutral, allowing it to stay active for longer times at even lower temperatures (see for example Alpar, Ankay, and Yazgan, 2001; Inutsuka and Sano, 2005; Alpar, Çalışkan, and Ertan, 2013). In general, the evolution of a fallback disk is not trivial to model, since it also depends on the complex interaction with the central compact source. However, if the disk becomes inactive and the accretion flow stops due to any of the mechanisms outlined above, this could explain the transition from the propeller back to the ejector phase.

In Figure 3.7, we show a possible evolutionary scenario for GLEAM-X J1627, GPM J1839–10 and PSR J0901-4046 that incorporates this drop in accretion rate. In particular, the top panels show the spin-period evolution. For all three sources we choose an initial magnetic field $B_0 = 4 \times 10^{14} \, \mathrm{G}$ and initial disk accretion rates of $\dot{M}_{\mathrm{d},0} = 10^{24} \, \mathrm{g \, s^{-1}}$ for GLEAM-X J1627 and $\dot{M}_{\mathrm{d},0} = 10^{23} \, \mathrm{g \, s^{-1}}$, for GPM J1839–10 and PSR J0901-4046 respectively. These values fall within the allowed parameter spaces that guarantee the sources to reach their spin period in less than $10^7 \, \mathrm{yr}$ (see Figure 3.6). Note that for PSR J0901-4046, we fix the value of B_0 to a higher value than the estimated current magnetic field strength in order to take into account the effects of magnetic field decay. To keep things simple, we set the disk accretion rate to zero as soon as the three neutron stars reach their observed spin periods of 1091 s, 1318 s and 75.9 s, respectively. This abrupt stop in the accretion process causes the sources to transition into the ejector phase and enables the restoration of the radio emission. From this point on, the spin periods remain effectively constant as the combination of their large values and magnetic field decay causes the dipolar losses to become negligible.

On the other hand, if accretion would persist, the neutron stars would continue
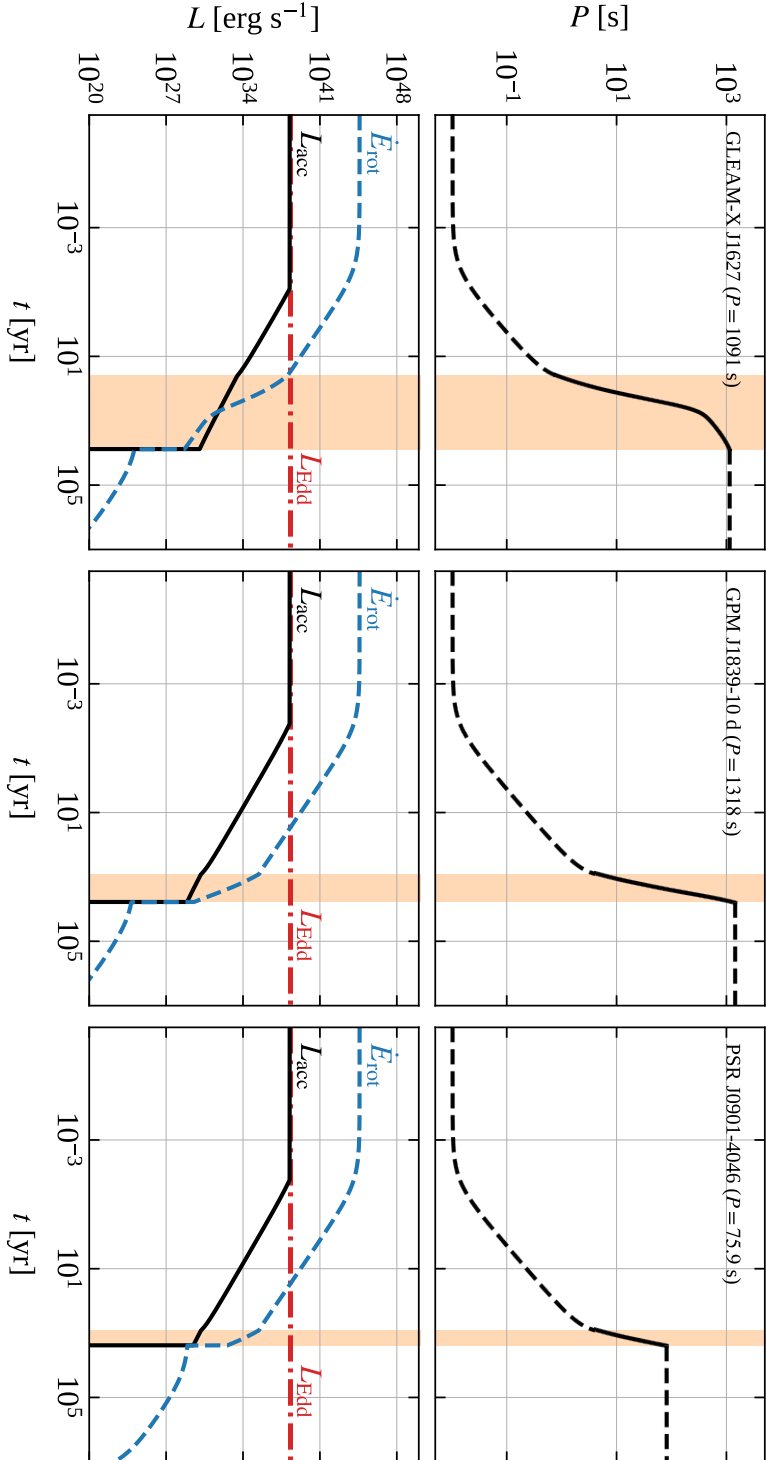
FIGURE 3.7.: Example of the time evolution of the spin period (top panels), disk luminosity and spin-down power (bottom panels) of GLEAM-X J1627 (left panels), GPM J1839-10 (central panels) and PSR J0901-4046 (right panels). For all sources we assume an initial magnetic field $B_0 = 4 \times 10^{14}$ G and an initial disk accretion rate of $\dot{M}_{d,0} = 10^{24}$ g s^{-1} for GLEAM-X J1627 and $\dot{M}_{d,0} = 10^{23}$ g s^{-1} for GPM J1839-10 and PSR J0901-4046, respectively. The orange-shaded region indicates the interval during which the three neutron stars spin down in the propeller phase. We assume that after the neutron stars have reached their observed periods of 1091 s, 1318 s and 75.9 s, respectively, an abrupt drop in accretion rate causes them to recover the ejector phase.

to evolve in a spin-equilibrium, unable to exit the propeller phase. This is possibly the evolutionary phase currently witnessed for 1E 161348-5055, the 2 kyr-old magnetar associated with the SNR RCW103. This source has shown magnetar-like outbursts (Rea et al., 2016; D'Aì et al., 2016) and an observed period of 6.67 hr but is detected only in the X-ray band. Its period evolution has been modeled by Ho and Andersson (2017) and Xu and Li (2019) using a similar model to our own. However, while Ho and Andersson (2017) assumed a constant accretion flow, in this work, we adopt the more realistic prescription of a time-varying accretion (similar to Xu and Li (2019)) as discussed in §3.3. In this scenario, the long period together with the young age of 1E 161348-5055 could be explained with a field of $\gtrsim 10^{15}$ G (in line with the results of Ho and Andersson, 2017; Xu and Li, 2019) and a relatively low initial disk accretion rate of $\sim 10^{21}$ g s^{-1} (see bottom right panel of Figure 3.4 as a reference). As 1E 161348-5055 is currently not observed in radio, and thus still in the propeller phase, we infer that the accretion disk is still active, which is compatible with the source's age of 2 kyr.

An association of long-period pulsars with supernova remnants could present additional proof for the fallback disk scenario. However, for GLEAM-X J1627, GPM J1839–10 and PSR J0901-4046 no clear evidence of SNR associations has been discovered so far. In general, SNRs are expected to have an observational lifetime of around 10^{4-5} yr (Braun, Goss, and Lyne, 1989; Leahy, Ranasinghe, and Gelowitz, 2020), which is comparable with our inferred timescales for the fallback disk to be active and accretion spinning down the three sources to their current periods. It is therefore possible that the associated SNRs are too faint to be detectable at the present time. Moreover, only around one-third of known young pulsars have a detected SNR remnant (see the ATNF Pulsar Catalog Manchester et al., 2005). The reasons for this are still a matter of debate and study, but the absence of detected SNRs for most neutron stars could simply indicate differences in their progenitors, supernova explosions and interstellar environments (see for example Gaensler and Johnston, 1995; Cui et al., 2021). We therefore do not consider the lack of SNR associations for GLEAM-X J1627, GPM J1839–10 and PSR J0901-4046 an issue for the validity of a fallback disk scenario.

3.9.2 Prediction for the X-ray and radio luminosities of long-period pulsars

In the bottom panels of Figure 3.7, we show the evolution of the disk luminosity and spin-down power for GLEAM-X J1627, GPM J1839–10 and PSR J0901-4046. The accretion luminosity is computed as in Equation (3.13), where we recall that $r_{\rm in} \simeq \min(r_{\rm m}, r_{\rm lc})$. At very early times, i.e. $\lesssim 1$ yr after the supernova, the accretion flow is limited by the Eddington limit, suggesting that the disk emits in X-rays at the Eddington luminosity. As the system evolves and the disk accretion rate starts to decrease as $t^{-\alpha}$, the inner disk radius increases following the evolution of the magnetospheric radius as $r_{\rm in} \sim r_{\rm m} \propto t^{2\alpha/7}$. As a consequence, the luminosity decreases roughly $\propto t^{-9\alpha/7}$. Once the disk becomes inactive or is completely consumed, the accretion rate is expected to vanish and the disk becomes undetectable in the X-rays. Therefore assuming the fallback scenario for long-period radio transients such as GLEAM-X J1627, GPM J1839–10 and PSR J0901-4046 which require ceasing of the accretion flow, we predict that X-ray observations should not be able to detect emission from a residual disk if present. The cold debris of inactive disks could instead be detectable in the infrared (Wang, Chakrabarty, and Kaplan, 2006; Posselt et al., 2018). However, if the central neutron stars are indeed young (around 10^5 yr) and
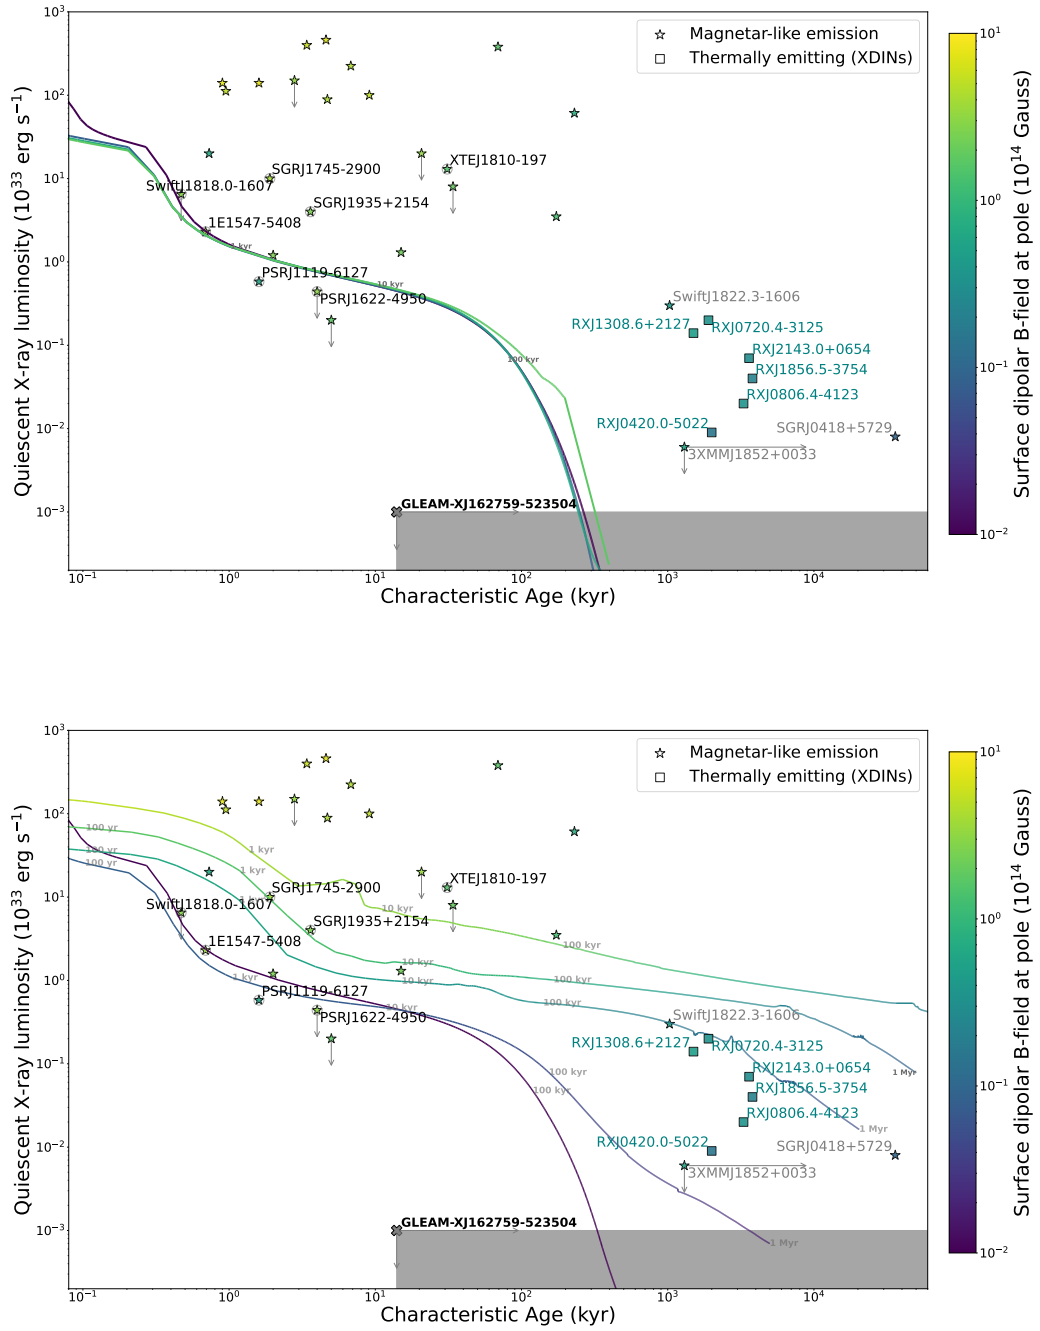
FIGURE 3.8: Evolutionary tracks for crust-dominated (*top panel*) and core-dominated (*bottom panel*) *B*-field configurations as a function of the characteristic age. Labels showing the real age are also marked along the curves in *grey*. Superimposed, are the X-ray luminosity of different pulsar classes, in particular radio-loud magnetars (*grey circles*), XDINSs (*light blue*) and low-field magnetars (*gray names*). The shaded region corresponds to the limits on the X-ray luminosity and the characteristic age derived for GLEAM-X J1627 (Figure taken from Rea et al., 2022).

have strong magnetic fields, they could emit thermal X-rays from their surfaces with luminosities up to $\sim 10^{31-35}\,\mathrm{erg\,s^{-1}}$ due to the dissipation of magnetic energy in the crust (Viganò et al., 2013).

However deep X-ray observations performed with the Chandra X-ray Observatory and XMM-Newton satellite have imposed upper limits of $\sim 10^{30}\,\mathrm{erg\,s^{-1}}$ (Rea et al., 2022) and $\sim 10^{33}\,\mathrm{erg\,s^{-1}}$ (Hurley-Walker et al., 2023) on the X-ray luminosity of GLEAM-X J1627 and GPM J1839–10, respectively. These limits challenge the magnetar interpretation especially for GLEAM-X J1627. Indeed comparing the X-ray luminosity of this source with cooling curves from magneto-thermal simulations (Viganò et al., 2021), the predicted age of GLEAM-X J1627 should be $> 1\,\mathrm{Myr}$ for any reasonable crustal-confined magnetic field ($B > 10^{13}\,\mathrm{G}$) (see upper panel in Figure 3.8). This age constraint is two orders of magnitude higher than that of typical radio-loud magnetars (which have estimated ages $< 20\,\mathrm{kyr}$). At this age, the bright radio bursts emitted by GLEAM-X J1627 would be unusual for such an old magnetar. On the other hand if the magnetar has a core-dominated magnetic field or has witnessed unusual fast cooling (both effects have never been unambiguously observed in a pulsar or a magnetar), the observed X-ray upper limits would be compatible with a younger age (see bottom panel in Figure 3.8).

In Figure 3.7 we also show the evolution of the spin-down power that is typically taken as the energy source for the radio emission of pulsars. In general, after these neutron stars have exited the propeller phase and recovered the ejector regime, we expect the spin-down to be caused by electromagnetic torques only so that $\dot{E}_{\rm rot} \propto B^2/P^4$ (see Equation (1.31) and (1.39)). Therefore, if we consider an upper limit for the magnetic field of around $10^{15}\,\mathrm{G}$ and a lower limit for the spin periods after the propeller regime of around $10\,\mathrm{s}$, we expect that the spin-down power has decayed to values $\lesssim 10^{33}\,\mathrm{erg\,s^{-1}}$ for long-period pulsar. This energy budget together with the magnetic energy stored in their strong fields is enough to power the observed radio emission and cause magnetar-like activity for neutron stars of this kind.

3.10 Summary

We have studied the spin evolution of young, isolated neutron stars under the influence of fallback accretion. We specifically focused on the fallback-disk scenario as a promising origin of long-period pulsars, a class of objects that recent radio surveys are starting to unveil. By solving the torque balance equation for a disk accreting neutron star, we demonstrate that the evolution of such an object can differ significantly from the standard dipole spin-down. In particular, we find that for a combination of high (but not extreme) magnetic field strengths and moderate fallback disk accretion rates in agreement with current core-collapse supernova simulations, neutron stars can enter the propeller phase during their evolution. This leads to effective spin-down and allows neutron stars to reach spin periods $\gg 10\,\mathrm{s}$ on time scales on the order of $\sim 10^{3-5}\,\mathrm{yr}$. Magnetic dipolar losses alone have problems explaining long spin periods and would require extreme conditions like strong and long-lasting magnetic fields potentially supported either by a core field component or other mechanisms such as the Hall attractor.

We have then interpreted the recently discovered objects GLEAM-X J1627, GPM J1839–10 and PSR J0901-4046, with rotation periods of 1091 s, 1318 s and 75.9 s, respectively, in light of this model, and showed that all three objects could be explained as highly magnetised neutron stars with a fallback disk accretion history. The possibility to reach such long spin periods in much less than $10^7\,\mathrm{yr}$ is crucial to maintain

the magnetic field and thus an energy reservoir to power their radio or X-ray activity. This is particularly important for GLEAM-X J1627, which was observed in outburst similar to other young radio-loud magnetars.

We showed that for newly born neutron stars with birth fields of $B_0 \sim 10^{14-15}$ G, a phase of fallback disk accretion with moderate initial accretion rates of 10^{22-27} g s^{-1}, could explain their detection as long-period radio or X-ray pulsars at relatively young ages ($\sim 10^{3-5}$ yr). On the other hand, in systems where the initial magnetic fields are lower ($\sim 10^{12-13}$ G), fallback disk accretion (even if present) is expected to have a negligible effect on the spin-period evolution. The majority of neutron stars will therefore primarily undergo standard dipolar electromagnetic spin-down and recover rotation periods below ~ 12 s. In our framework, we therefore naturally recover the pulsar population which is observed to spin in the range ~ 0.002–12 s (see Figure 1.12). Note that the recently discovered radio pulsars PSR J1903+0433 (Han et al., 2021) and PSR J0250+5854 (Tan et al., 2018) with periods of 14 s and 23 s, respectively, could be easily accommodated within our fallback accretion scenario. However, both sources can in principle also be explained within the standard evolutionary scenario provided that crustal field decay is very weak (essentially requiring the absence of a highly resistive pasta layer; see Pons, Viganò, and Rea, 2013).

We also mentioned the X-ray emitting magnetar at the centre of the 2 kyr-old SNR RCW103, which requires (ongoing) fallback to explain its 6.67 hr period and radio-quiet nature. Finally note that classical magnetars with periods $\lesssim 12$ s would correspond to those systems where only strong fields are present, but fallback disk accretion does not take place or is inefficient because of a low $\dot{M}_{d,0}$.

In conclusion, fallback disk accretion after the supernova explosions of massive stars is expected to affect the evolution of newly born neutron stars. Depending on the relative intensities of the initial pulsar magnetic field and accretion rates, this scenario could represent an important ingredient to explain the connection between different neutron star classes and specifically shed light on the nature of the long-period radio sources recently discovered.

Chapter 4

Population study of long-period pulsars in the neutron star and white dwarf rotating dipole scenarios

4.1 Introduction

As we have discussed in Chapter 3, the discovery of ultra-long periodic coherent radio emitters challenges our current knowledge of neutron-star emission and evolution. Some of these new sources seem to be extreme neutron-star pulsars (e.g., the 76-s source PSR J0901-4046; Caleb et al. 2022), while the interpretation of others is still uncertain (e.g., the 18-minute source GLEAM-X J1627 Hurley-Walker et al. 2022a, and the 21-minute GPM J1839–10 Hurley-Walker et al. 2023). The periodic radio emission features of GLEAM-X J1627 and GPM J1839–10 are similar to other radio magnetars, and their long periodicity can be explained through past supernova fallback accretion (see Chapter 3, e.g., Alpar, Ankay, and Yazgan, 2001; Chatterjee, Hernquist, and Narayan, 2000; Ertan et al., 2009; Tong et al., 2016; Ronchi et al., 2022). X-ray observations with the Chandra X-ray Observatory and XMM-Newton satellite allowed to derive deep X-ray luminosity limits of $\lesssim 10^{30}\,\mathrm{erg\,s^{-1}}$ and $\lesssim 10^{33}\,\mathrm{erg\,s^{-1}}$ for GLEAM-X J1627 and GPM J1839–10 , respectively that challenge the magnetar interpretation (Rea et al., 2022; Hurley-Walker et al., 2023). By considering these upper limits, magneto-thermal evolution models (Viganò et al., 2013; Viganò et al., 2021) predict cooling ages of these two new sources that exceed 1 Myr, far beyond any currently observed magnetar (Rea et al., 2022).

In contrast, slow spin periods are common in magnetic white dwarfs (Ferrario and Wickramasinghe, 2005; Ferrario, Wickramasinghe, and Kawka, 2020). Although isolated magnetic white dwarfs exhibit magnetic-field strengths between 10^6 to $10^9\,\mathrm{G}$ (Ferrario, de Martino, and Gänsicke, 2015; Ferrario, Wickramasinghe, and Kawka, 2020), lower than neutron star B-fields spanning 10^8 to $10^{15}\,\mathrm{G}$ (see also Figure 1.12, magnetic white dwarfs have also been proposed to emit spin-down-driven radio emission similar to neutron stars (Zhang and Gil, 2005). To date, two radio-emitting white dwarfs have been detected, the binary systems AR Sco ($P \sim 1.95\,\mathrm{min}$ in a 3.5 hr orbit; Marsh et al. 2016) and J1912−4410 ($P \sim 5.3\,\mathrm{min}$ in a 4 hr orbit; Pelisoli et al. 2023). The radio emission of both systems is partly compatible with dipolar spin-down (Geng, Zhang, and Huang, 2016; Buckley et al., 2017; du Plessis et al., 2019), but also has a significant component resulting from the intrabinary shock with the wind of the companion star. The lack of an optical/infrared counterpart to GLEAM-X J1627 at the estimated distance of 1.3 kpc rules out a similar binary

system for this source (Rea et al., 2022). However, it does not exclude lower mass companions or the possibility of a relatively cold isolated magnetic white dwarf.

In this chapter based on the work in Rea, Hurley-Walker, Pardo-Araujo, Ronchi, Graber, Coti Zelati, De Martino, Bahramian, McSweeney, Galvin, Hyman, and Dall'Ora (2023), we study GLEAM-X J1627 and GPM J1839–10 in the context of the radio emission expected from the spin-down of an isolated neutron star and white dwarf by means of death-line analyses (Section 4.3) and population-synthesis simulations (Section 4.4). In the following section I will review briefly our current knowledge on magnetic white dwarfs.

4.2 Magnetic white dwarfs

To date we have detected around 600 isolated white dwarfs exhibiting magnetic fields in the range $10^3 - 10^9$ G and another 200 in binary systems (see Ferrario, de Martino, and Gänsicke, 2015; Ferrario, Wickramasinghe, and Kawka, 2020, for a review). As white-dwarf atmospheres are rich in hydrogen and helium, the Zeeman splitting of spectral lines provides a direct measure of the magnetic field strength on the surface of these stars. Circular polarisation in the spectrum, caused by the presence of a strong magnetic field, is another indicator of the magnetism in white dwarfs (Ferrario, de Martino, and Gänsicke, 2015; Ferrario, Wickramasinghe, and Kawka, 2020; Caiazzo et al., 2021). Studies of the local distribution of white dwarfs have revealed that between 10% and 20% of the white dwarf population exhibits relevant magnetic fields (see, e.g., Holberg et al., 2016; Bagnulo and Landstreet, 2021).

The origin of magnetism in white dwarfs is not well understood. Different processes can play a role in generating strong magnetic fields during the evolution of white dwarfs and their progenitors also depending on their physical properties. As for the neutron star case (see Section 1.5) one hypothesis is the fossil origin scenario. In this case the magnetic field is inherited from the progenitor star and is amplified due to flux conservation during the contraction to a white dwarf. This require a progenitor already having a strong magnetic field. It is commonly believed that the progenitors of magnetic white dwarfs are Ap and Bp-type stars which usually possess a strong magnetic field up to tens of kiloGauss (Ferrario and Wickramasinghe, 2005). However, given the incidence of magnetism in the white-dwarf population Kawka and Vennes (2004) showed that the local density of Ap and Bp stars is not sufficient to account for the entire magnetic white dwarf population. Therefore, it has also been suggested that dynamo processes during binary interactions and stellar mergers can play a role in shaping the magnetism of these objects (Tout et al., 2008). An alternative theory is that magnetic fields in white dwarfs can be produced and amplified due to convective processes during the crystallisation of the liquid core at times of around 1 Gyr after the white dwarf formation (Isern et al., 2017). This hypothesis is supported also by observational evidence of an increasing occurrence of magnetism in older white dwarfs (Bagnulo and Landstreet, 2021). Observations have also shown that magnetic white dwarfs tend to be on average more massive than the non-magnetic ones with an excess mass of around 0.07 $M_\odot$ (Liebert, 1988; Kepler et al., 2013; Bagnulo and Landstreet, 2021). This supports the theory that magnetic white dwarfs originate from massive A and B-type progenitors, close to the limit of around 8 $M_\odot$ above which neutron stars are created. Also, the merger scenario can be a viable explanation since after the merger a more massive outcome is expected (García-Berro et al., 2012).
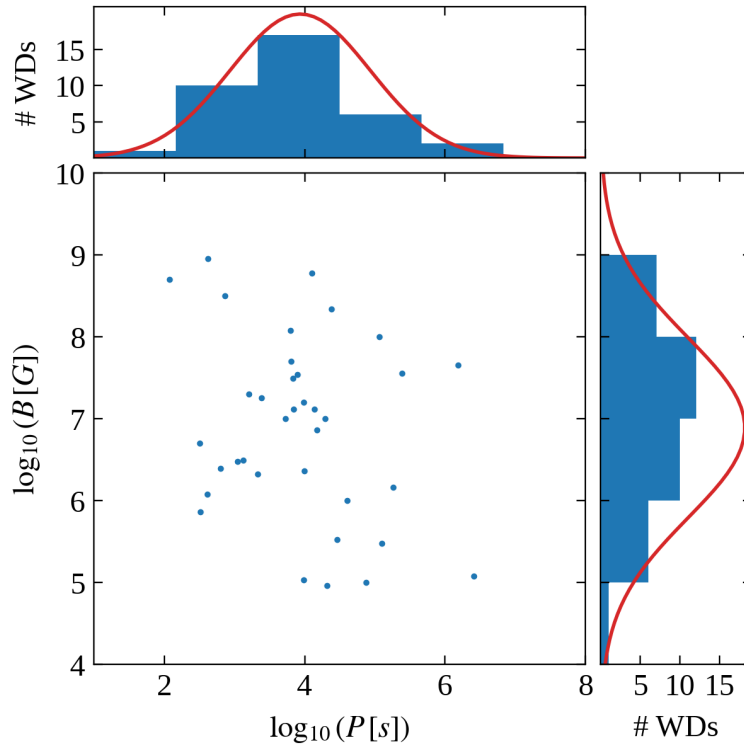
FIGURE 4.1: Relation between the logarithm of magnetic field strengths and spin periods of the currently known magnetic white dwarfs with a reliable measurement of these two quantities (see Ferrario, Wickramasinghe, and Kawka (2020)). On the top and right panels we report the corresponding histograms for the two quantities with the corresponding fit with Gaussian distributions (see text for more details).

In general, white dwarf spin periods are not easily determined. For magnetic white dwarfs, the magnetic activity and the presence of magnetic spots on the surface can induce spectroscopic, polarimetric and photometric variability that together with asteroseismology techniques can be used to estimate the spin period (see Wickramasinghe and Ferrario, 2000; Ferrario, de Martino, and Gänsicke, 2015; Ferrario, Wickramasinghe, and Kawka, 2020, and references therein). The measured spin periods are in the range of a few hundred seconds up to decades or even centuries and show that white dwarfs are slow rotators. Such a large range of spin periods can not be explained with angular momentum conservation during the contraction phase and is likely evidence for different physical processes playing a role during the formation and evolution of white dwarfs. For example, efficient angular momentum transfer from the core to the envelope and large-scale angular momentum loss during post-main-sequence evolution can contribute to slowing down the rotation. Moreover, magnetic fields are likely to lead to increased braking of the stellar core as it approaches the white dwarf state. Binary interaction and accretion phases during the progenitor evolution and mergers also play a role in shaping the final spin period (Wickramasinghe and Ferrario, 2000). Figure 4.1 shows the magnetic field versus rotation period of all currently known magnetic white dwarfs. We note that there is no significant correlation between field strength and spin period.

4.3 Death valleys for neutron-star and white-dwarf radio pulsars

Radio emission from rotating magnetospheres is usually explained as a result of pair production just above the polar caps (Ruderman and Sutherland, 1975). However, as we have seen in Section 1.9.3, for certain limiting periods, magnetic-field strengths and geometries, radio pulsars can no longer produce pairs, and radio emission ceases.

The parameter space in the $P - B$ plane (or equivalently $P - \dot{P}$ plane as B and $\dot{P}$ are related via Equation (1.46)) below which radio emission is quenched is called the "death valley" (Chen and Ruderman, 1993; Zhang, Harding, and Muslimov, 2000, see also Section 1.9.3). This death valley encompasses a large variety of death lines depending on the magnetic-field configuration (e.g., dipolar, multi-polar, twisted), the nature of the seed gamma-ray photons for pair production (i.e., curvature or inverse Compton photons), the pulsar obliquity, the stellar radius and moment of inertia (see Suvorov and Melatos 2023 for a death-valley discussion for long-period pulsars). To date, these death-line models have been applied exclusively to neutron star pulsars because, until very recently, no white dwarf pulsed radio emission had been detected. However, magnetic white dwarfs might not be unlike neutron star pulsars in generating coherent radio emission through magnetospheric spin-down losses (Zhang and Gil, 2005), albeit with different stellar radii, masses and magnetic fields.

Figure 4.2 shows death valleys for neutron star and white dwarf pulsars as red and blue-shaded regions, respectively. Their boundaries are marked by two death-line extremes (Chen and Ruderman, 1993), representing the broadest range of B-field configurations (see Section 1.9.3). I report them here for convenience:

- a pure dipole configuration (see Equation (1.71)):

$$B_{\mathrm{p}} \gtrsim 2.0 \times 10^{12}\,\mathrm{G} \left(\frac{R}{10^6\,\mathrm{cm}} \right)^{-19/8} \left(\frac{P}{1\,\mathrm{s}} \right)^{15/8} ; \qquad (4.1)$$

- an extremely twisted, multipolar magnetic field located in a small spot above the polar cap (see Equation (1.72)):

$$B_{\mathrm{p}} \gtrsim 8.3 \times 10^{10}\,\mathrm{G}\, b^{-1/4} \left(\frac{R}{10^6\,\mathrm{cm}} \right)^{-2} \left(\frac{P}{1\,\mathrm{s}} \right)^{3/2} . \qquad (4.2)$$

   where b is the ratio between the spot's B-field and the dipolar strength B_{p} (here we assume an extreme value of $b = 10$ as in Section 1.9.3).

These death lines can be also represented in the $P - \dot{P}$ diagram by considering Equation (1.46) as an estimate of the magnetic field strength at the pole and making explicit the dependence on the mass and radius:

$$B_{\mathrm{p}} = \left(\frac{3c^3}{5\pi^2} \right)^{1/2} M^{1/2} R^{-2} P^{1/2} \dot{P}^{1/2} \qquad (4.3)$$

$$\simeq 5.7 \times 10^{19}\,\mathrm{G} \left(\frac{M}{1M_\odot} \right)^{1/2} \left(\frac{R}{10^6\,\mathrm{cm}} \right)^{-2} P^{1/2} \dot{P}^{1/2},$$
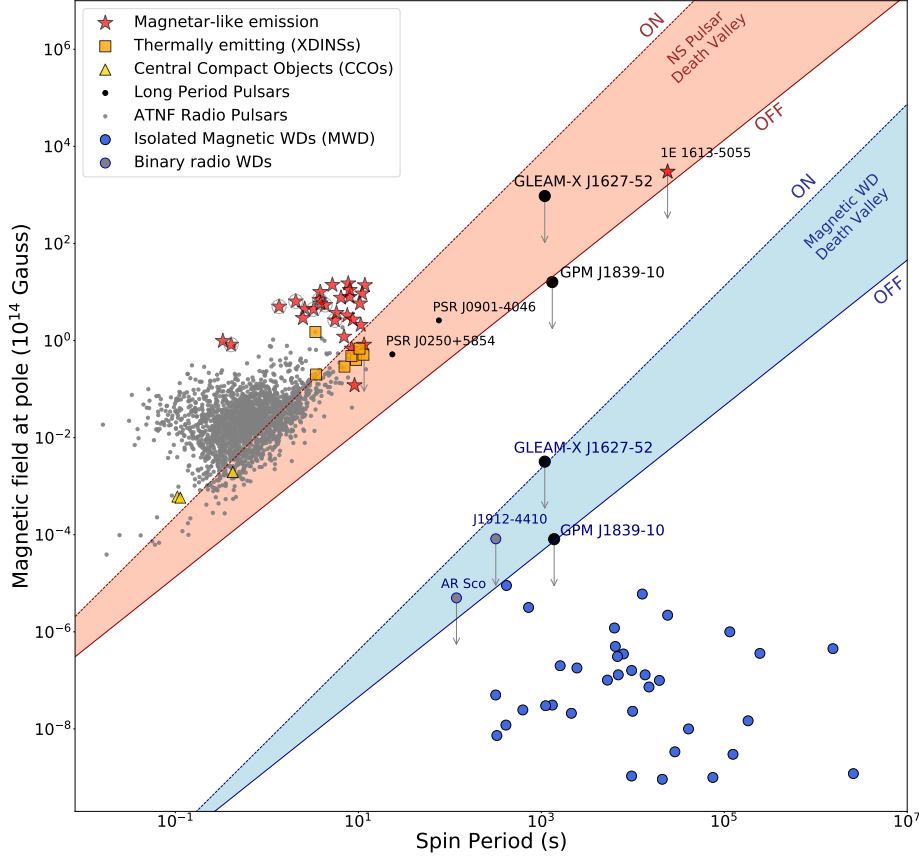
FIGURE 4.2: Surface dipolar magnetic field, *B*, against spin period, *P*, for observed isolated neutron stars and magnetic white dwarfs. GPM J1839–10 and GLEAM-X J1627 are interpreted as isolated neutron stars or white dwarfs. Arrows represent upper *B*-field limits. We show isolated ATNF radio pulsars (Manchester et al., 2005) (gray dots), pulsars with magnetar-like X-ray emission (red stars; gray circles highlight the radio magnetars), including the long-period magnetar 1E 161348-5055 (De Luca et al., 2006; Rea et al., 2016; D'Aì et al., 2016), X-ray Dim Isolated neutron stars (XDINSs; orange squares) and Central Compact Objects (CCOs; gold triangles) (Olausen and Kaspi, 2014; Coti Zelati et al., 2018). Other long-period radio pulsars are reported as black circles (Tan et al., 2018; Caleb et al., 2022). Isolated magnetic white dwarfs are represented by blue dots (Ferrario, Wickramasinghe, and Kawka, 2020; Caiazzo et al., 2021; Buckley et al., 2017). Gray dots show upper *B*-field limits for the binary white dwarfs AR Sco (Buckley et al., 2017) and J1912-4410 (Pelisoli et al., 2023) computed from Equation 4.6. Dashed (solid) lines correspond to theoretical death lines for a pure dipole (extremely twisted multipole) configuration. Red and blue shaded regions indicate neutron star and white dwarf death valleys, respectively.

where we assumed a moment of inertia $I = 2/5MR^2$. By substituting Equation 4.3 into Equations 4.1 and 4.2 we obtain the following corresponding death lines:

$$\dot{P} \gtrsim 1.2 \times 10^{-15}\,\mathrm{s\,s^{-1}} \left(\frac{M}{1M_\odot}\right)^{-1} \left(\frac{R}{10^6\,\mathrm{cm}}\right)^{-3/4} \left(\frac{P}{1\,\mathrm{s}}\right)^{11/4}; \qquad (4.4)$$

$$\dot{P} \gtrsim 2.1 \times 10^{-18}\,\mathrm{s\,s^{-1}} b^{1/2} \left(\frac{M}{1M_\odot}\right)^{-1} \left(\frac{P}{1\,\mathrm{s}}\right)^2. \qquad (4.5)$$

Note that the second death line turns out to be independent of the radius of the star.

Although these death lines rely on simplifications compared to more recent works (e.g., Zhang, Harding, and Muslimov, 2000), our focus is on the extremes of the death valley. All newer models incorporating more detailed physics fall within the shaded region for any reasonable neutron star or white dwarf parameters. For Figure 4.2, we use a fiducial neutron star radius of $R_\mathrm{NS} = 11$ km, in line with recent measurements (Riley et al., 2021; Raaijmakers et al., 2021, see Section 1.3). For white dwarfs, we assume $R_\mathrm{WD} = 6000$ km consistent with the Hamada-Salpeter relation (Hamada and Salpeter, 1961) and measurements of isolated magnetic white dwarfs (Ferrario, de Martino, and Gänsicke, 2015; Ferrario, Wickramasinghe, and Kawka, 2020).

We also add observed neutron-star pulsars and magnetic white dwarfs to the $P - B$ plane in Figure 4.2. For observed neutron star pulsars, we derive surface magnetic fields at their poles using P and $\dot{P}$ measurements, employing the dipolar loss formula (Equation (1.46)) assuming $M_\mathrm{NS} = 1.4M_\odot$ and $R_\mathrm{NS} = 11$ km. B-fields at the poles of isolated white dwarfs are obtained from observations (i.e., Zeeman splitting of spectral lines; Ferrario, de Martino, and Gänsicke 2015; Ferrario, Wickramasinghe, and Kawka 2020; Caiazzo et al. 2021). For the radio pulsating white dwarfs AR Sco and J1912−4410, we estimate upper B-field limits assuming the emission to result from dipolar losses alone (see Buckley et al., 2017). By combining Equations (1.31) and (1.44) and assuming an inclination angle $\chi = \pi/2$, we find the expression:

$$B_\mathrm{P} \sim \left(\frac{3c^3}{8\pi^4}\right)^{1/2} R_\mathrm{WD}^{-3} P^2 L^{1/2}$$

$$\sim 4.7 \times 10^8\,\mathrm{G} \left(\frac{R_\mathrm{WD}}{6000\,\mathrm{km}}\right)^{-3} \left(\frac{P}{100\,\mathrm{s}}\right)^2 \left(\frac{L}{10^{33}\,\mathrm{erg\,s^{-1}}}\right)^{1/2}. \qquad (4.6)$$

For AR Sco and J1912−4410-like systems, Equation (4.6) represents an upper limit on the B-field since, besides magnetic dipolar losses, the bulk of the spin-down power is likely lost by a combination of different effects: magnetohydrodynamic (MHD) interactions of the white dwarf's magnetic field with the secondary star, outflows of relativistic charged particles from the magnetic white dwarf and winds from the companion star (Marsh et al., 2016; Buckley et al., 2017).

Finally, we also show the upper limits on the surface dipolar B-fields of the two long-period radio sources GLEAM-X J1627 and GPM J1839–10 in Figure 4.2.

4.4 Population synthesis for neutron-star and white-dwarf radio pulsars

We simulate isolated neutron star and white dwarf populations using the framework of Chapter 6 (see also Ronchi et al. (2021) and Chapter 5) with model parameters adjusted for each object type. Initially, we randomly sample the logarithm of the birth periods and magnetic fields from normal distributions, and the inclination angle between the magnetic and the rotational axis from a uniform distribution in spherical coordinates. Assuming that neutron stars and white dwarfs spin down due to magnetospheric torques, we then evolve their periods, P, and inclination angles, χ, over time by solving the coupled differential equations (Spitkovsky, 2006; Philippov, Tchekhovskoy, and Li, 2014, see also Section 1.8):

$$\dot{P} = \frac{\pi^2}{c^3} \frac{B^2 R^6}{IP} \left(\kappa_0 + \kappa_1 \sin^2 \chi \right), \tag{4.7}$$

$$\dot{\chi} = -\frac{\pi^2}{c^3} \frac{B^2 R^6}{IP^2} \left(\kappa_2 \sin \chi \cos \chi \right), \tag{4.8}$$

where we assume for simplicity $I = 2/5(MR^2)$ and $\kappa_0 \simeq \kappa_1 \simeq \kappa_2 \simeq 1$ for pulsars surrounded by magnetospheres (see discussion in Section 1.8). Finally, we determine the number of stars that point towards the Earth by assuming a random direction for the line of sight and employing a prescription for the aperture of the radio beam as outlined below (see Chapter 6 for more details).

To compare the impact of various initial model assumptions on the final spin-period distributions, we carry out the population simulations summarised in Table 4.1 and Figures 4.3–4.5. Specifically, in Table 4.1, to help the reader, we count the objects falling within the period ranges $10 - 10^2$ s, $10^2 - 10^3$ s, and $10^3 - 10^5$ s. We then distinguish objects intercepting our line of sight and those with $\dot{E} > 10^{27}$ erg s^{-1} (see Figures 4.3–4.5 for the exact $\dot{E}$ and P distributions). The latter limit has no intrinsic meaning, but was chosen as a reference to show how many sources would have sufficient rotational power to support GPM J1839–10's radio luminosity.

4.4.1 Neutron star population synthesis

We simulate 10^7 neutron stars with random ages sampled from a uniform distribution up to a maximum age of 10^9 yrs. This translates to a birth rate of one neutron star per century, consistent with the Galactic core-collapse supernova rate (Rozwadowska, Vissani, and Cappellaro, 2021). To assign each neutron star a birth field, we then sample the logarithm of the field (in Gauss) from a normal distribution with mean $\mu_{\log B} = 13.25$ and a standard deviation of $\sigma_{\log B} = 0.75$ (see, e.g., Gullón et al., 2014; Gullón et al., 2015; Cieślar, Bulik, and Osłowski, 2020). Unless stated otherwise, we adopt $M_{NS} = 1.4\,M_\odot$ and $R_{NS} = 11$ km. We sample the logarithm of the initial period from a normal distribution with mean $\mu_{\log P}$ = -0.6 (corresponding to 0.25 s) and standard deviation $\sigma_{\log P} = 0.3$ (Popov et al., 2010; Gullón et al., 2014; Xu et al., 2023).

We further incorporate magnetic-field decay due to Ohmic dissipation and the Hall effect through magneto-thermal evolution curves from Viganò et al. (2013) and Viganò et al. (2021) (see also Chapter 6 for more details) and assume a radio beam angular aperture $\propto P^{-1/2}$ (see Equation (1.57) and Lorimer and Kramer, 2012) where we assume an emission radius $r_{em} = 300$ km (Johnston and Karastergiou, 2019).
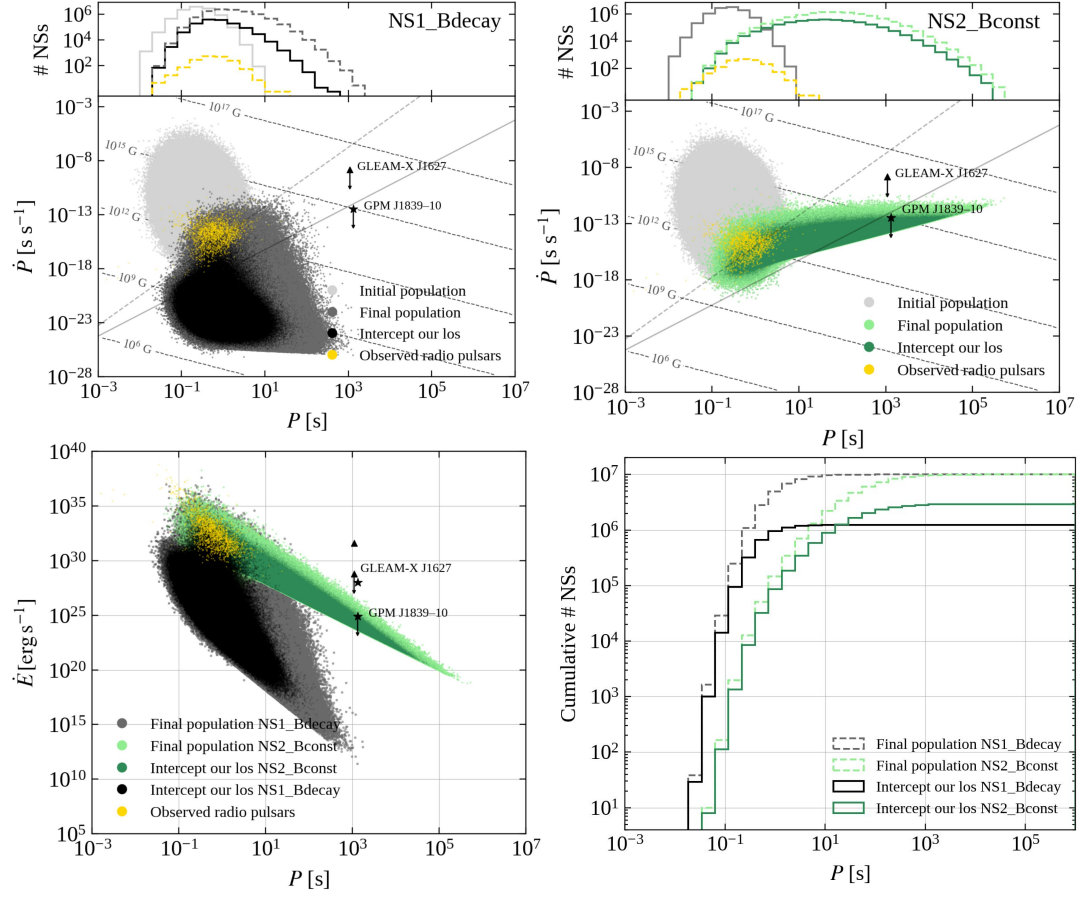
FIGURE 4.3: Population-synthesis results for models `NS1_Bdecay` (black) and `NS2_Bconst` (green). Top panels show $P - \dot{P}$ diagrams for both simulations, respectively. Light grey dots represent initial neutron star populations, dark grey and light green final populations. The subsets of objects intercepting our line of sight (los) are shown in black and dark green. Lines of constant B-fields and $\dot{E}$ and the death lines defined in Equations 4.4 and 4.5 are indicated for reference. Histograms above the $P - \dot{P}$ diagrams represent the corresponding period distributions. The bottom left panel shows $\dot{E}$ versus spin period for the evolved populations, while the bottom right panel highlights the cumulative period distributions. With the black triangle and black star we also report the two sources GLEAM-X J1627 and GPM J1839–10. In the $\dot{E}$ versus spin period diagram we report both the upper limit on the $\dot{E}$ and the estimated radio luminosity for the two sources. Yellow dots highlight the observed isolated pulsar population given in the ATNF pulsar catalogue (Manchester et al., 2005).
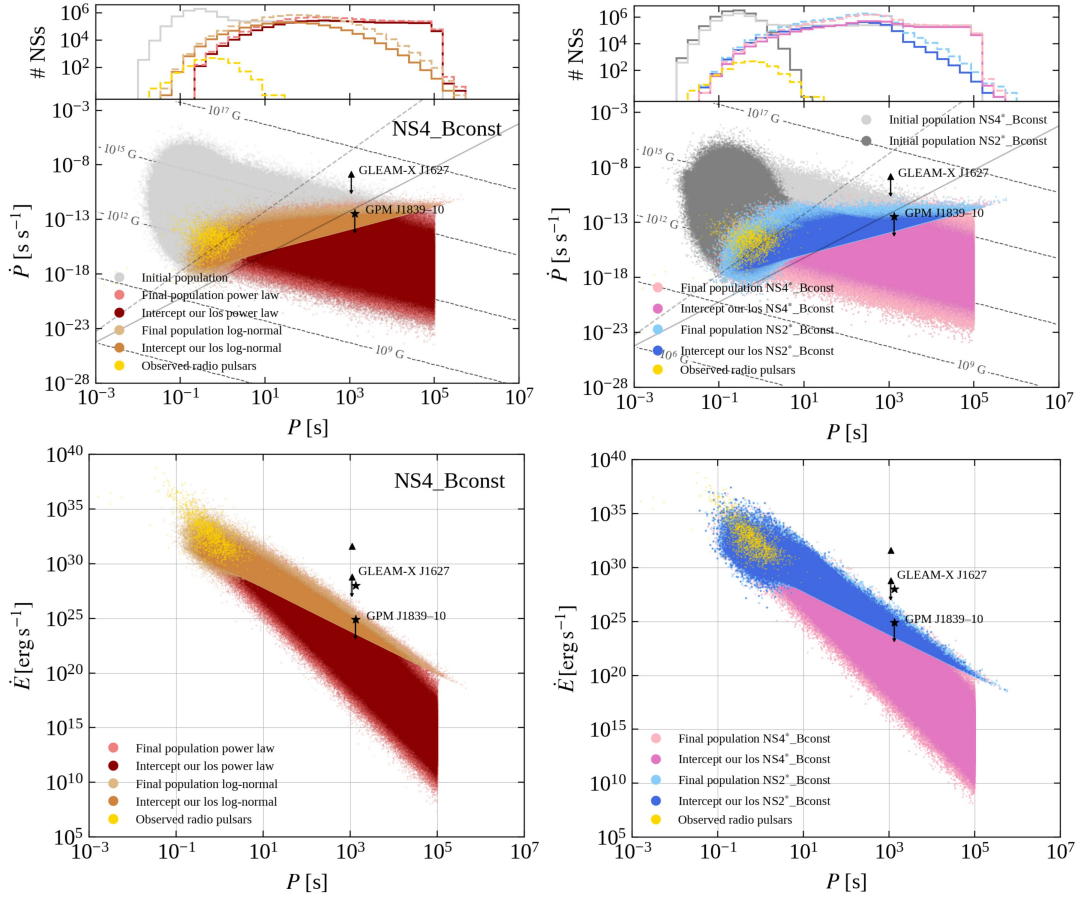
FIGURE 4.4: Population-synthesis results for models `NS4_Bconst` (orange/red), `NS2*_Bconst` (blue), and `NS4*_Bconst` (pink). Panels, lines and yellow dots are equivalent to those in Figure 4.3. In both left panels, evolved objects sampled from the log-normal (power-law) contribution to the initial period distribution are shown in orange (red). Right panels show evolved populations of models `NS4*_Bconst` and `NS2*_Bconst` based on a bimodal *B*-field distribution. Across all panels, light shades depict evolved populations, while objects intercepting our lines of sight are shown in dark shades.

Model `NS1_Bdecay` serves as a reference with standard population assumptions (Figure 4.3 top-left). These are typical initial parameters compatible with the current observed pulsar population (yellow dots in Figures 4.3 and 4.4). However, they cannot predict any long-period pulsars.

We continue with investigating more extreme scenarios, focusing first on an evolution without field decay. Strong fields could be maintained over long timescales if electric currents are predominantly present in the neutron star core (e.g., Viganò et al., 2021). Consequently, neutron stars experience a more pronounced spin-down, reaching longer periods. For model `NS1_Bconst`, we thus repeat the set-up of `NS1_Bdecay` but with constant B-field at the very limit of what is physically viable. However, adding the constant-B assumption is insufficient to slow down the population substantially (see Table 4.1). For subsequent models, we continue with the extreme constant B-field case to explore the impact of other assumptions.

In model `NS2_Bconst`, we also relax the standard beaming assumption, adjusting the radio beam angular aperture to obtain a duty cycle of roughly 20% independent from the spin period P, in line with observations of GLEAM-X J1627 and GPM J1839–10. To do so, we consider an orthogonal rotator and compute the beam aperture that would result in a duty cycle of 20% ($\delta = 0.2$) for this configuration. This implies that the angular aperture of the beam has to be $\rho_{em} = \delta\pi = 36°$ (see also Section 1.9. This adjustment results in an increase of the number of pulsars crossing our line of sight especially in the long period tale (see Figure 4.3 top-right and bottom panels). For the remaining simulations, we thus maintain this prescription of the beaming unless stated otherwise.

Next, we explore different initial spin-period distributions, mimicking a possible interaction with initial fallback accretion (see, e.g., Alpar, Ankay, and Yazgan, 2001; Ertan et al., 2009; Tong et al., 2016; Ronchi et al., 2022). For models `NS3_Bconst` to `NS6_Bconst`, we add a power law with an arbitrary cut-off at a period of 10^5 s to the aforementioned log-normal distribution of the observed pulsar population. We specifically consider a power law, as the spin-down is likely determined by different fallback accretion rates. Note that the cut-off does not affect our final results, but reflects the maximum spin reachable by fallback accretion (see Figures 3.4 and 3.5 and Ronchi et al., 2022). We arbitrarily assume that both distributions are equally normalised, sampling 50% of neutron stars from either distribution, maintaining a birth rate of 1 neutron star per century. This prescription is still consistent with the log-normal population resulting in the observed radio pulsars (Gullón et al., 2015, see also yellow dots in Figures 4.3 and 4.4). For models `NS3_Bconst` and `NS4_Bconst` (see Figure 4.4 left panels), we assume a corresponding power-law index of -3 and -1, respectively. `NS5_Bconst` investigates a duty cycle of 10% ($\delta = 0.1$), i.e., $\rho_{em} = \delta\pi = 18°$, while for `NS6_Bconst`, we explore the effect of the assumed mass, setting $M_{NS} = 2M_\odot$.

Since stronger magnetic fields enhance the spin-down, we also investigate the effect of a bimodal B-field distribution (four models denoted with an asterisk ($*$) in Table 4.1). In particular, besides the log-normal distribution, we consider that 50% of neutron stars are formed with a strong field uniformly distributed in $\log B \in [13.5, 14.5]$ following Gullón et al. (2015). `NS2*_Bdecay` and `NS2*_Bconst` consider only the log-normal for the initial period distribution and a decaying and constant magnetic field, respectively, while for `NS4*_Bdecay` and `NS4*_Bconst`, we explore the log-normal plus power law for the initial period (see Figure 4.4 right panels).
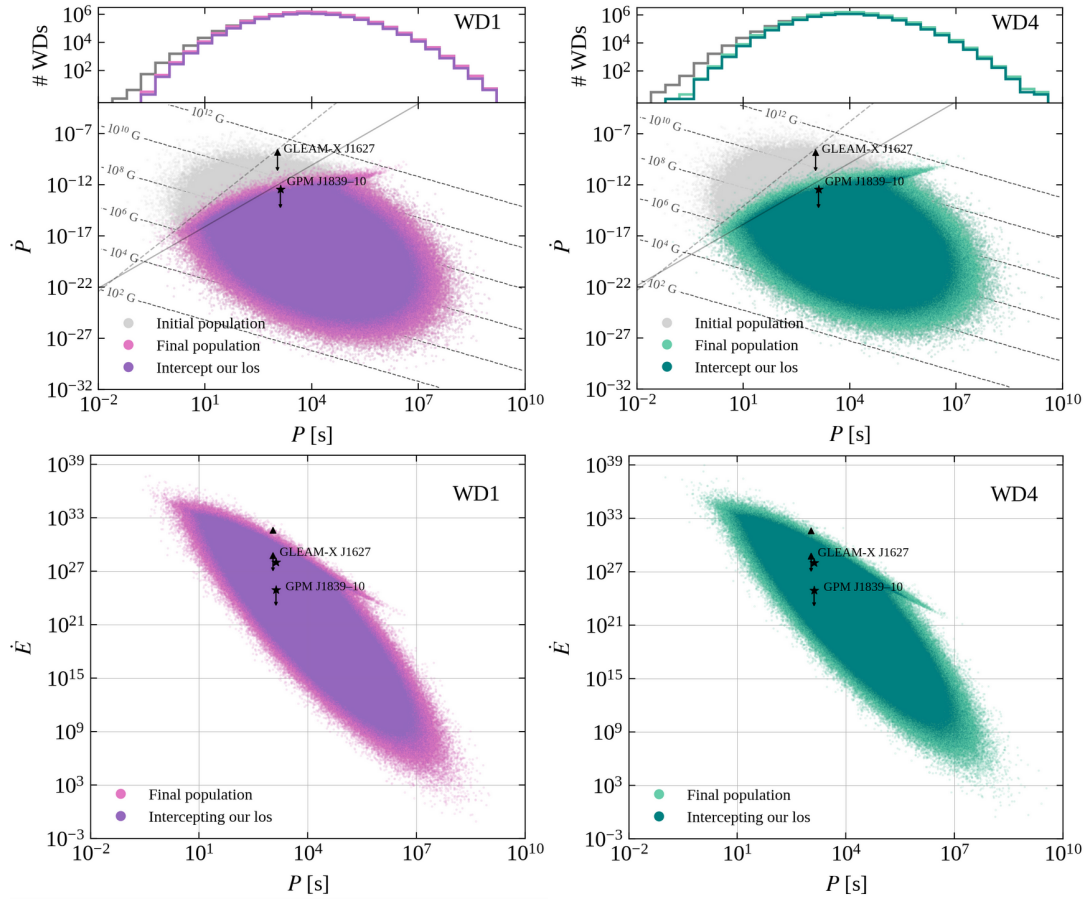
FIGURE 4.5: Population-synthesis results for models WD1 (purple) and WD4 (seagreen). Panels and lines are similar to Figure 4.3 and Figure 4.4. Gray dots represent initial white dwarf populations, pink and light seagreen final populations. The subsets of objects intercepting our line of sight are shown in purple and dark seagreen.

4.4.2 White dwarf population synthesis

Magnetic white dwarfs spin down slower than neutron stars due to larger moments of inertia, larger spin periods and lower *B*-fields (see Equation 1.42). Moreover, magnetic fields of magnetic white dwarfs do not exhibit relevant magnetic-field decay due to longer Ohmic dissipation timescales (e.g., Cumming, 2002) and can be taken as constant (Ferrario, Wickramasinghe, and Kawka, 2020). Considering Equation (1.14) we obtain a rough estimate of the Ohmic timescale in white dwarfs:

$$\tau_{\mathrm{Ohm}} \simeq 4.4 \times 10^9 \, \mathrm{yr} \left(\frac{\sigma}{10^{21} \, \mathrm{s}^{-1}} \right) \left(\frac{L}{1000 \, \mathrm{km}} \right)^2, \tag{4.9}$$

where we consider typical values of the conductivity from Cumming (2002). Consequently, current isolated white dwarf periods and magnetic-field strengths closely reflect those at birth.

To model these birth distributions, we consider a sample of 37 magnetic white dwarfs with reliable spin-period and magnetic-field measurements (Ferrario, Wickramasinghe, and Kawka, 2020). We fit Gaussian functions to the distributions of the logarithm of the periods and *B*-fields, deriving a mean of $\mu_{\log P} = 3.94$ and standard deviation of $\sigma_{\log P} = 1.0$, and $\mu_{\log B} = 6.91$ and $\sigma_{\log B} = 1.09$, respectively (see Figure 4.1). The white dwarf birth rate in the Milky Way has been predicted to be around 100 per century assuming a Galaxy radius of 20 kpc and a Galactic disk height of 0.5 kpc (Holberg et al., 2016). For our population synthesis, we then simulate 10^8 magnetic white dwarfs with ages drawn from a uniform distribution up to a maximum of 10^9 yr, consistent with a birth rate of 10 per century assuming 10% of the white dwarfs being magnetic (see, e.g., Holberg et al. 2016 but also Bagnulo and Landstreet 2021 who recently found 22%). We then assign initial *P* and *B* values from our fitted distributions. Results of four simulation configurations are summarised in Table 4.1 and Figure 4.5.

We further assume a white dwarf radio beam angular aperture independent of *P*. For models WD1 and WD2, we adjust our approach to obtain a 20% and 10% duty cycle, corresponding to beam apertures $\rho_{\mathrm{em}} = 18°$ and $36°$, respectively (see Section 4.4.1). For WD3 and WD4, we set the beaming as in WD1 but vary mass and radius. In particular, using the Hamada-Salpeter mass-radius relation for He white dwarfs (Hamada and Salpeter, 1961), we consider $M_{\mathrm{WD}} = 1.2 \, M_\odot$ and $R_{\mathrm{WD}} = 4000 \, \mathrm{km}$ for a high-mass white dwarf in WD3 and $M_{\mathrm{WD}} = 0.6 \, M_\odot$ with $R_{\mathrm{WD}} = 9000 \, \mathrm{km}$ for a low-mass white dwarf in WD4.

TABLE 4.1: Population-synthesis results for isolated neutron stars and magnetic white dwarfs. All models were evolved for 10^9 yr assuming a birth rate of 1 neutron star and 10 magnetic white dwarf per century. Unless specified otherwise, neutron stars have $1.4 M_\odot$ and 11 km radii, while magnetic white dwarfs have $1 M_\odot$ and 6000 km radii. Numbers in parentheses denote the assumed power-law (PL) slope. Numbers in brackets are the sub-samples of simulated neutron stars that cross our line of sight [l.o.s.] for their respective beaming. Note that we normalise the neutron star and white dwarf numbers to 10^3.

Model	Initial P	Beaming	Total N/10^3 [l.o.s.]			Total N/10^3 with $\dot{E} > 10^{27}$ erg s^{-1} [l.o.s.]		
			$P = 10^{1-2}$ s	$P = 10^{2-3}$ s	$P = 10^{3-5}$ s	$P = 10^{1-2}$ s	$P = 10^{2-3}$ s	$P = 10^{3-5}$ s
			NEUTRON STARS					
NS1_Bdecay	Log-N	pulsars	629.51 [9.22]	6.69 [0.02]	0.01 [0.00]	2.49 [0.05]	0.01 [0.00]	0.00 [0.00]
NS1_Bconst	Log-N	pulsars	4560.95 [31.15]	3357.42 [3.70]	574.38 [0.09]	2797.74 [24.68]	34.62 [0.06]	0.00 [0.00]
NS2_Bconst	Log-N	20%	4564.26 [1360.90]	3355.02 [799.39]	573.61 [120.38]	2799.49 [876.59]	34.70 [8.75]	0.00 [0.00]
NS3_Bconst	Log-N+PL(-3)	20%	4568.02 [1417.86]	3364.06 [818.03]	573.68 [122.11]	2799.46 [915.60]	34.78 [9.09]	0.00 [0.00]
NS4_Bconst	Log-N+PL(-1)	20%	3425.36 [1302.58]	3344.05 [1341.96]	2291.68 [1475.93]	1920.09 [706.18]	26.38 [8.52]	0.00 [0.00]
NS5_Bconst	Log-N+PL(-1)	10%	3425.41 [603.07]	3345.82 [646.41]	2289.66 [838.73]	1920.40 [320.11]	26.44 [3.62]	0.00 [0.00]
NS6_Bconst^O	Log-N+PL(-1)	20%	3551.79 [1364.31]	3158.14 [1298.75]	2200.36 [1453.03]	2269.94 [837.21]	37.55 [12.22]	0.00 [0.00]
NS2_Bdecay*	Log-N	20%	827.50 [367.95]	3.37 [0.93]	0.00 [0.00]	3.75 [1.74]	0.00 [0.00]	0.00 [0.00]
NS2_Bconst*	Log-N	20%	2621.41 [765.23]	5849.46 [1342.22]	774.43 [168.48]	1656.60 [504.92]	77.12 [19.16]	0.00 [0.00]
NS4_Bdecay*	Log-N+PL(-1)	20%	1511.96 [977.49]	895.09 [665.54]	1786.98 [1330.37]	3.92 [2.27]	0.01 [0.00]	0.00 [0.00]
NS4_Bconst*	Log-N+PL(-1)	20%	1942.87 [717.60]	5077.88 [1733.97]	2508.41 [1565.64]	1136.06 [406.37]	58.33 [18.82]	0.00 [0.00]
			WHITE DWARFS					
WD1	Log-N	20%	1929.55 [1406.56]	14485.85 [10649.08]	69050.05 [51234.73]	1823.72 [1328.54]	7956.49 [5803.80]	3880.60 [2729.91]
WD2	Log-N	10%	1928.76 [817.16]	14484.90 [6225.08]	69042.10 [30034.14]	1822.90 [771.63]	7955.59 [3386.88]	3883.30 [1577.85]
WD3^O	Log-N	20%	2147.21 [1576.08]	14705.23 [10871.34]	68592.09 [51016.48]	1894.94 [1389.56]	5687.17 [4172.56]	1664.71 [1176.52]
WD4^O	Log-N	20%	1579.74 [1141.63]	13837.27 [10076.93]	70003.37 [51617.09]	1542.90 [1114.49]	9672.34 [6989.46]	7872.76 [5454.37]

* Bimodal magnetic field distribution (see text).

^{O} NS6_Bconst assumes a $2M_\odot$ neutron star mass, WD3 a mass of $1.2M_\odot$ and 4000 km radius, and WD4 a mass of $0.6M_\odot$ and 9000 km radius.

4.5   Results and discussion

Figure 4.2, provides an overview of long period sources in the classical scenario of neutron star pulsar radio emission (see section 1.9.3) based on magnetospheric pair production. While this scenario can in principle accommodate GLEAM-X J1627, it cannot account for GPM J1839–10 as the source sits below even the most extreme death line. However, note that both objects have radio luminosities exceeding their $\dot{E}$s by 2-3 orders of magnitudes (see Hurley-Walker et al., 2022a; Hurley-Walker et al., 2023, and Section 3.6, and 3.7) implying that dipolar spin-down alone cannot explain the observed luminosities. Hence, the emission scenario for long-period emitters is necessarily more complex than for normal radio pulsars, possibly resembling radio magnetars, where the electromagnetic emission can be powered by the energy stored in the strong magnetic field (see also Rea et al., 2022, and Section 1.13).

Figure 4.2 also highlights that a similar mechanism in magnetic white dwarfs could in principle contribute to the radio emission of GLEAM-X J1627, AR Sco and J1912−4410 in the white-dwarf scenario. However, GPM J1839–10's bright radio emission cannot be easily reconciled even in the isolated magnetic white dwarf case. For white dwarfs in binary systems such as AR Sco and J1912−4410, the interactions of the companion's wind with the white dwarf's magnetosphere can replenish the magnetosphere with plasma and enhance the radio emission (Geng, Zhang, and Huang, 2016). Deep optical and IR observations ruled out main sequence stars for GLEAM-X J1627 (Rea et al., 2022), but deeper observations are needed to exclude any binarity. On the other hand, GPM J1839–10's limits (Hurley-Walker et al., 2023) cannot provide strong constraints given its larger distance.

Our neutron star population synthesis models summarised in Table 4.1 show that a large population of long-period radio emitters can not be easily explained as neutron star pulsars. Neither standard population assumptions nor the most extreme scenarios invoking no field decay (Figure 4.3), initial slow-down via fallback accretion, 20% duty cycles or stronger birth fields (Figure 4.4) result in sufficiently energetic neutron star pulsars with periods > 1000 s pointing towards the Earth (irrespective of mass). A difference by a factor of a few in the neutron star birth rate does not alter this conclusion.

On the other hand, white dwarf population synthesis highlights that long-period magnetic white dwarfs are more common than neutron star pulsars. This is primarily because they are born with longer spin periods. This combined with their lower magnetic field and larger moment of inertia reduce the effectiveness of the dipolar spin-down process. Only white dwarfs with small spin periods and large magnetic field experience appreciable spin down (see Figure 4.5). Therefore the abundance of slow rotators (Table 4.1) in the white-dwarf scenario mainly reflects their birth properties. However, Figure 4.2 shows that the known sample of isolated magnetic white dwarfs are not expected to emit coherent radio emission via standard pair production, being all located below the most extreme death line.

4.6   Conclusions

In this chapter, we studied long-period pulsars in the rotating neutron-star and white-dwarf dipole scenario, one of the most likely interpretations given their coherent and highly polarised emission. We also perform population synthesis simulations to try to asses the likelihood of detecting long-period sources originating from these scenarios.

We find that the classical particle acceleration mechanism for rotating dipoles (see Section 1.9.3) fails to provide a satisfactory explanation for the radio emission of GPM J1839–10 in either the neutron star or white dwarf scenario. In contrast, all observed isolated magnetic white dwarfs with measured *B*-fields fall below the most extreme death lines, possibly explaining their radio non-detection. The radio emission observed from the binary white dwarfs AR Sco and J1912−4410 might be enhanced by the presence of their companion star within the white dwarf pulsar light cylinder. However, for GLEAM-X J1627 optical and IR observations could rule out main sequence companion stars (Rea et al., 2022). For GPM J1839–10 a similar constrain was not possible given the larger distance (Hurley-Walker et al., 2023).

Moreover, in the neutron star scenario, we do not expect a large population of ultra-long period pulsars to possess enough energy to power the observed coherent emission under any (physically motivated or extreme) assumptions. While many more slow white dwarf pulsars can be expected possessing enough rotational energy, we however still lack a mechanism to explain the bright radio emission. Therefore, if GLEAM-X J1627 and GPM J1839–10 are confirmed as isolated neutron star or white dwarf pulsars, this would call for a revision of our understanding of radio emission from dipolar magnetospheres. Corroborating the neutron star scenario would further require a significant re-examination of our understanding of initial neutron star parameters (birth rates, magnetic-field distribution, etc.) and their evolution at the population level.

Chapter 5

Analyzing the Galactic pulsar dynamics with machine learning

5.1 Introduction

Neutron stars have been observed to travel through the Galaxy with typical velocities of around several hundreds of kilometers per second, reaching more than a thousand kilometers per second in some extreme cases (Chatterjee et al., 2005; Hobbs et al., 2005; Hui and Becker, 2006; Pavan et al., 2014). Accurate information on neutron star positions and velocities in the Milky Way usually comes from radio timing and interferometric observations (see Chapters 8 and 9 in Lorimer and Kramer, 2012; Liu et al., 2020, and references therein) or high spatial resolution X-ray observations with, e.g., the Chandra X-ray observatory (Motch et al., 2009). These observations provide measurements of the pulsars' angular positions in the sky and their proper motions projected onto the celestial sphere. In some cases, the radio pulse dispersion measure (DM) or the X-ray absorption density (N_H) together with Galactic free electron-density and hydrogen-density models (Balucinska-Church and McCammon, 1992; Taylor and Cordes, 1993; Cordes and Lazio, 2002; Yao, Manchester, and Wang, 2017) can also yield a rough distance estimate. Moreover, in a few cases, a parallax measurement (Deller et al., 2009; Matthews et al., 2016; Wang et al., 2017; Deller et al., 2019) or the presence of a Supernova Remnant (SNR) (Yao, Manchester, and Wang, 2017) might provide better distance measurements.

Such high proper velocities of the neutron star population as a whole exceed those of their progenitors (typically massive OB stars) (see Hansen and Phinney, 1997; Lai, Chernoff, and Cordes, 2001, and references therein), and cannot be explained by the neutron stars' motion in the Galactic gravitational potential alone. The mechanisms providing such high velocities are still unclear but are likely related to the underlying supernova explosion. One possibility is that the central core of an exploding star receives a kick due to an asymmetric ejection of material from the star's outer layers; a direct result of momentum conservation (Shklovskii, 1970; Dewey and Cordes, 1987; Mandel and Müller, 2020). Additionally, the anisotropic emission of neutrinos has been suggested to impart kicks on compact remnants (Bisnovatyi-Kogan, 1993; Fryer and Kusenko, 2006; Tamborra et al., 2014; Nagakura, Sumiyoshi, and Yamada, 2019).

However, constraining the neutron stars' natal kick-velocity distribution from current observational data is not straightforward. Most pulsars, especially those with very high velocities, have moved far away from their birth places, and their proper motions have been modified by the Galactic gravitational potential. Thus, the current velocity of a pulsar may differ substantially from its velocity at birth. Knowing the exact pulsar age and its current 3D spatial velocity, we are in principle able to recover the initial conditions by integrating the pulsar's orbit back in time.

However, in general we lack information about the pulsar's line-of-sight velocity, and accurate knowledge about its age, since the characteristic age estimated from the pulsar period and its derivative can differ significantly from the true age (see e.g. Kaspi et al., 2001; Viganò et al., 2013, and Section 1.8). Furthermore, estimates of pulsar distances have typically large associated errors due to uncertainties in the underlying density models used to convert pulsar DM or N_{H} into distance estimates (Lorimer et al., 2006; He, Ng, and Kaspi, 2013; Deller et al., 2019).

Reconstruction of the three-dimensional initial position and velocity distribution of pulsars, and comparison with the observed Galactic neutron star population is therefore a complicated task that requires careful simulations as well as detailed estimates of the observational biases of multi-band surveys. Several studies have performed statistical and population synthesis analyses to recover the distributions of important neutron star parameters from the observed population (see e.g. Arzoumanian, Chernoff, and Cordes, 2002; Brisken et al., 2003; Hobbs et al., 2005; Faucher-Giguère and Kaspi, 2006; Gullón et al., 2014; Verbunt, Igoshev, and Cator, 2017; Cieślar, Bulik, and Osłowski, 2020). While these models are broadly able to explain the observational data, high degrees of degeneracy between the different input parameters make it difficult to exactly pin down the distributions that control the pulsars' birth properties, such as their natal kick velocities. Nonetheless, disentangling the birth properties of the isolated neutron star population in our Galaxy is crucial as it has important implications for several lines of research, including formation mechanisms of these compact stars, the evolution of massive stars, as well as extreme events such as Gamma-Ray Bursts (GRBs), Fast Radio Bursts (FRBs) and peculiar types of supernovae.

In this chapter, based on the work published in Ronchi, Graber, Garcia-Garcia, Rea, and Pons (2021), we focus on characterising initial pulsar properties using machine learning techniques (see Chapter 2). I will describe the technical aspects related to these efforts and show that a machine-learning framework can be used to estimate parameters with high accuracy. For this feasibility study, we restrict ourselves to a simplified approach, where selection effects and observational biases are neglected and reduced physical models are sufficient. In particular, we focus on the dynamical properties of the pulsar population and explore the possibility of inferring the parameters that control a given Galactic pulsar kick-velocity and scale-height distribution at birth (the two quantities that largely control the spatial distribution of pulsars in the Milky Way) through neural networks. For this purpose, we implement a basic population synthesis code in `Python` and simulate the dynamical evolution of a synthetic population of isolated neutron stars for a variety of different birth-position and natal-kick distributions. These evolved mock populations are then fed into a suitably structured machine-learning pipeline with the aim of recovering the underlying parameters controlling the distributions. We show that this procedure is successful at estimating birth characteristics. Additionally, we link our framework to the observed sample of pulsars with measured proper motion in a phenomenological way and discuss implications for future pulsar survey efforts, e.g., with the Square Kilometer Array (SKA).

In Section 5.2, I describe the methods used to simulate and evolve a mock neutron star population in time. Section 5.3 contains a description of the machine-learning framework, including the generation of our datasets (Section 5.3.1), the employed network architectures (Section 5.3.2) as well as details of the training process (Section 5.3.3). In Section 5.4, I present our experiments, which are discussed in detail and connected with observational data in Section 5.5. Finally, I provide a summary and outlook in Section 5.6.

5.2 Population synthesis: dynamics

A widely used approach to investigate the properties of the observed neutron star population is through population synthesis (see e.g. Narayan and Ostriker, 1990; Faucher-Giguère and Kaspi, 2006; Gonthier et al., 2007; Kiel et al., 2008; Kiel and Hurley, 2009; Osłowski et al., 2011; Levin et al., 2013; Gullón et al., 2014; Bates et al., 2014; Cieślar, Bulik, and Osłowski, 2020, and Section 1.15). These frameworks aim to simulate the evolution of a population of neutron stars from birth until today. The resulting mock population is then compared with the real observed population in order to constrain and validate the physical model assumptions that entered the simulation. In particular, the population synthesis approach relies on assumptions about the distributions of the birth properties of the mock neutron stars, and typically takes advantage of Monte–Carlo methods to construct the initial parameters of each simulated star. Starting from these initial conditions, the mock population is then evolved over time according to some evolutionary prescriptions, and eventually contrasted with real data. For the development of our population synthesis framework, we largely follow Faucher-Giguère and Kaspi (2006), Gullón et al. (2014) and Cieślar, Bulik, and Osłowski (2020). In this chapter we will focus on the dynamical properties while in the next chapter we will extend the analysis to the magneto-rotational properties of neutron stars. The necessary ingredients are briefly summarised in the following.

5.2.1 Age

The age, t_{age}, of each neutron star is randomly drawn from a uniform probability distribution between 1 and 10^7 yr. By choosing a uniform distribution, we assume that the birth rate of neutron stars is constant in the chosen time range. For all simulations of the synthetic neutron star population, we choose an average neutron star birth rate of 1 neutron star per century, compatible with the core-collapse supernova rate in the Galaxy (Rozwadowska, Vissani, and Cappellaro, 2021, see also Section 1.14). This yields a total of 10^5 simulated neutron stars for each synthetic population, whose evolution we can compute within reasonable timescales.

5.2.2 Birth position

To define the initial positions, we use both a Cartesian reference frame (x, y, z) and a cylindrical reference frame (r, ϕ, z), whose origins are located at the centre of the Galaxy. Here, r represents the distance in kiloparsec from the Galactic centre, ϕ is the azimuthal angle in radians and z is the distance from the Galactic plane. The two coordinate systems are related by the transformation:

$$\begin{cases} x = r\cos\phi, \\ y = r\sin\phi, \\ z = z. \end{cases} \tag{5.1}$$

We assume that the Sun is located at the coordinates $x = 0\,\mathrm{kpc}$, $y = r_\odot$, $z = z_\odot$, where $r_\odot = 8.3\,\mathrm{kpc}$ and $z_\odot = 0.02\,\mathrm{kpc}$ (see Pichardo et al., 2012, and references therein). We calculate the initial position at birth of each neutron star in both cylindrical and Cartesian galactocentric reference frames. To do so, we execute the following steps:

i) First, we draw a random distance r from the Galactic centre for each neutron star ranging between 10^{-4} and $20\,$kpc according to a pulsar radial density distribution $P(r)$. In particular, we follow the Milky Way's pulsar surface density $\rho(r)$ defined by Equation (15) in Yusifov and Küçük ([2004](#)) to determine the probability density function for the radial distance, i.e.:

$$P(r) = 2\pi r\rho(r), \tag{5.2}$$

with

$$\rho(r) = A \left(\frac{r + r_1}{r'_\odot + r_1} \right)^a \exp\left[-b \left(\frac{r - r'_\odot}{r'_\odot + r_1} \right) \right], \tag{5.3}$$

where $A = 37.6\,$kpc^{-2}, $a = 1.64$, $b = 4.0$, $r_1 = 0.55\,$kpc and $r'_\odot = 8.5\,$kpc is the Sun's distance from the Galactic centre projected on the Galactic disk. Although different from the $r_\odot$ value assumed above, we keep $r'_\odot = 8.5\,$kpc in this parametrisation in order to be consistent with the results of Yusifov and Küçük ([2004](#)). We note that this is the distribution for evolved pulsars rather than that of their progenitors and Yusifov and Küçük ([2004](#)) find small discrepancies between this distribution and that of OB stars. However, Faucher-Giguère and Kaspi ([2006](#)) show that the evolved pulsar population is well described by birth positions drawn from the density distribution in Equation ([5.3](#)) and argue that differences fall within the current uncertainties of pulsar distance measurements. Given the lack of a more realistic description, we therefore adopt the above prescription.

ii) Neutron stars are born mainly within the Galactic spiral arms, as these regions are rich in massive OB stars (Chen et al., [2019](#)). We implement a model for the galactic spiral structure that includes four arms with a logarithmic shape function, which gives the azimuthal coordinate, ϕ, as a function of the distance from the Galactic centre:

$$\phi(r) = k \ln\left(\frac{r}{r_0} \right) + \phi_0. \tag{5.4}$$

Our values of the model parameters, i.e., the winding constant k, the inner radius r_0 and the inner angle ϕ_0 are reported in Table [5.1](#) and evaluated from Table 1 in Yao, Manchester, and Wang ([2017](#)) in order to match the same functional form as defined in Equation ([5.4](#)). For our analysis, we follow Faucher-Giguère and Kaspi ([2006](#)) and do not include the Local arm, whose density is much smaller than that of the four major arms (Cordes and Lazio, [2002](#); Yao, Manchester, and Wang, [2017](#)). For each star, we then randomly select one of the four spiral arms with equal probability, and evaluate the angular coordinate, ϕ, for its given r according to Equation ([5.4](#)).

The spiral pattern of the Galaxy is not static and as a first approximation can be considered as a rigid structure which rotates with an approximated period $T = 250\,$Myr (Vallée, [2017](#); Skowron et al., [2019](#)). Knowing the age of an object and assuming a rotational angular velocity of $\Omega = 2\pi/T$ for the spiral structure, we can derive the angular position at birth of each neutron star. Note that the Galaxy rotates in the clockwise direction, i.e., toward decreasing ϕ angles.

After obtaining the corresponding angular coordinate for each neutron star

TABLE 5.1: Parameters of the Milky Way spiral arm structure: the winding constant k, the inner radius r_0 and the inner angle ϕ_0 (adapted from Table 1 in Yao, Manchester, and Wang (2017); see Section 5.2.2 for more details).

Arm number	Name	k [rad]	r_0 [kpc]	ϕ_0 [rad]
1	Norma	4.95	3.35	0.77
2	Carina-Sagittarius	5.46	3.56	3.82
3	Perseus	5.77	3.71	2.09
4	Crux-Scutum	5.37	3.67	5.76

birth position, we add noise to both coordinates r and ϕ to smear out the distribution and avoid artificial features near the Galactic centre. For this purpose, we add a correction $\phi_{\mathrm{corr}} = \phi_{\mathrm{rand}} \exp(-0.35r/\mathrm{kpc})$ to the ϕ coordinate, where ϕ_{rand} is randomly drawn from a uniform distribution in the interval $[0, 2\pi)$, and to the r coordinate a correction r_{corr} randomly drawn from a normal distribution centred at 0 with standard deviation $\sigma = 0.07r$. Although this prescription was introduced by Faucher-Giguère and Kaspi (2006) (see their Section 3.2.1) in a somewhat arbitrary manner, the resulting stellar distribution broadly agrees with that observed for very young high-mass stars as shown in Reid et al. (2019).

Then the birth position in polar coordinates of each neutron star is given by $(r + r_{\mathrm{corr}}, \phi(r) + \Omega t_{\mathrm{age}} + \phi_{\mathrm{corr}})$ with units [kpc, rad].

iii) To determine the height z in kiloparsec from the Galactic plane of each neutron star, we adopt an exponential disk model as given by Wainscoat et al. (1992). It is shaped by the characteristic scale-height parameter h_{c}:

$$P(z) = \frac{1}{h_{\mathrm{c}}} \exp\left(-\frac{|z|}{h_{\mathrm{c}}}\right). \tag{5.5}$$

For our machine-learning experiments, we will vary the scale height in the range $[0.02, 2]$ kpc to simulate neutron star populations with different spread in Galactic height. This range encompasses the value $h_{\mathrm{c}} = 0.18$ kpc, which was adopted by Gullón et al. (2014) to match radio pulsar observations, and which is also compatible with the population of young massive stars in the Galactic disk (Li et al., 2019). We will consider $h_{\mathrm{c}} = 0.18$ kpc below, whenever a fiducial scale height is required for our synthetic pulsar population. The coordinate z of each neutron star is then randomly drawn according to this height probability distribution in a range of 10^{-4} to 5 kpc. We choose a maximal distance of 5 kpc from the Galactic plane to model a fixed galactic volume for all of our simulation runs, while also ensuring sufficient resolution for the objects in the galactic disc for those models with small scale heights. Subsequently, for each star we randomly choose if z is positive or negative determining in this way a position above or below the Galactic plane.

TABLE 5.2: Parameters of the Milky Way gravitational potential taken from Table 1 in Marchetti, Rossi, and Brown (2019); see also Bovy (2015).

Component	Parameters
nucleus (n)	$M_\mathrm{n} = 1.71 \times 10^9 M_\odot$
	$R_\mathrm{n} = 0.07$ kpc
bulge (b)	$M_\mathrm{b} = 5.0 \times 10^9 M_\odot$
	$R_\mathrm{b} = 1.0$ kpc
disk (d)	$M_\mathrm{d} = 6.8 \times 10^{10} M_\odot$
	$a_\mathrm{d} = 3.00$ kpc
	$b_\mathrm{d} = 0.28$ kpc
halo (h)	$M_\mathrm{h} = 5.4 \times 10^{11} M_\odot$
	$R_\mathrm{h} = 15.62$ kpc

5.2.3 Initial velocity

We assume that the initial velocity of the neutron stars in the Galaxy is given by two contributions: the progenitor velocity in the Galactic gravitational potential and a kick speed imparted onto the neutron stars as a result of the supernova explosion (see also Section 1.4). We consider a progenitor circular orbital speed given by the following relation:

$$v_\mathrm{orb} = \sqrt{r \frac{\partial \Phi_\mathrm{MW}(r,z)}{\partial r}}, \tag{5.6}$$

where Φ_MW is the Milky Way gravitational potential discussed below. We assume that each neutron star has an initial kick velocity v_k, whose 3D magnitude v_k is randomly drawn from a Maxwell distribution, shaped by the dispersion parameter σ_k:

$$P(v_\mathrm{k}) = \sqrt{\frac{2}{\pi}} \frac{v_\mathrm{k}^2}{\sigma_\mathrm{k}^3} \exp\left(-\frac{v_\mathrm{k}^2}{\sigma_\mathrm{k}^2}\right). \tag{5.7}$$

For our machine-learning purposes, we will vary σ_k in the range $[1, 700]$ km s^{-1} and randomly draw 3D velocity magnitudes from the resulting distribution in the range $[0, 2500]$ km s^{-1}. This spread allows us to easily accommodate the fastest observed neutron stars whose velocities have been estimated to surpass 1000 km s^{-1} (see for example Chatterjee et al., 2005; Hui and Becker, 2006; Pavan et al., 2014). Based on pulsar timing measurements, Hobbs et al. (2005) have suggested that $\sigma_\mathrm{k} = 265$ km s^{-1} provides a viable explanation for the proper motions of observed neutron stars. We will use this as a fiducial value below. For a given kick-velocity magnitude, we then associate a random direction to this velocity in order to evaluate the three components ($v_{\mathrm{k},r}$, $v_{\mathrm{k},\phi}$, $v_{\mathrm{k},z}$) in galactocentric cylindrical coordinates. Therefore, the three components of the total initial velocity of each neutron star in the galactocentric reference frame are computed as ($v_{\mathrm{k},r}$, $v_\mathrm{orb} + v_{\mathrm{k},\phi}$, $v_{\mathrm{k},z}$).

5.2.4 Galactic potential

As typical for spiral galaxies like the Milky Way, we assume an axisymmetric Galactic potential Φ_{MW} (Carlberg and Innanen, 1987; Bovy, 2015) that does not incorporate the impact of the spiral arms themselves. We specifically consider a four-component Galactic potential model consisting of a nucleus (Φ_n), a bulge (Φ_b), a disk (Φ_d) and a halo (Φ_h) as discussed in Marchetti, Rossi, and Brown (2019):

$$\Phi_{MW} = \Phi_n + \Phi_b + \Phi_d + \Phi_h. \tag{5.8}$$

The nucleus and the bulge are described by a spherical Hernquist potential (Hernquist, 1990):

$$\Phi_n = -\frac{GM_n}{R_n + R}, \tag{5.9}$$

$$\Phi_b = -\frac{GM_b}{R_b + R}. \tag{5.10}$$

where the coordinate $R = \sqrt{r^2 + z^2}$ is the distance from the Galactic centre in spherical coordinates. The disk has a Miyamoto-Nagai disk potential (Miyamoto and Nagai, 1975):

$$\Phi_d = -\frac{GM_d}{\sqrt{K^2 + r^2}}, \tag{5.11}$$

where $K = a_d + \sqrt{z^2 + b_d^2}$ is the shape parameter with a_d as the scale length and b_d the scale height of the disk. The halo has a Navarro-Frenk-White potential (Navarro, Frenk, and White, 1996):

$$\Phi_h = -\frac{GM_h}{R} \ln\left(1 + \frac{R}{R_h}\right). \tag{5.12}$$

The parameters of this model are reported in Table 5.2 and were derived by Bovy (2015) through a fit of the mass profile of the Milky Way. We assume that these contributions to the galactic potential are stationary in time, i.e., they do not evolve over the time span we consider for the dynamical evolution.

5.2.5 Dynamical evolution

Given the initial conditions defined above, i.e., the initial position, initial velocity and the Galactic gravitational potential, we can solve the equations of motion to determine the neutron stars' dynamical evolution. The system of dynamical equations that requires solving to determine the orbits of the neutron stars in the galactic potential is given by the Newtonian equations of motion: $\ddot{r} = -\nabla\Phi_{MW}$. In cylindrical galactocentric coordinates the three components of this vector equation take the

form:

$$
\begin{cases}
\ddot{r} - r\dot{\phi}^2 = -\dfrac{\partial \Phi_{\mathrm{MW}}}{\partial r}, \\[2ex]
2\dot{r}\dot{\phi} + r\ddot{\phi} = -\dfrac{1}{r}\dfrac{\partial \Phi_{\mathrm{MW}}}{\partial \phi}, \\[2ex]
\ddot{z} = -\dfrac{\partial \Phi_{\mathrm{MW}}}{\partial z}.
\end{cases}
\tag{5.13}
$$

For each neutron star, we numerically integrate the above equations in time from $t = 0$ yr to $t = t_{\mathrm{age}}$ using a discrete time step. We use the Python package `scipy.integrate.odeint` and to speed up the computational time also employ the module `jit` ("just in time") from the `Numba` package (https://numba.pydata.org/, Lam, Pitrou, and Seibert 2015).[1] To asses the performance of our integration method, we check that both the total energy (i.e., the potential plus kinetic energy) and the z-component of the total angular momentum of all the stars in our simulation are conserved. For simplicity, we assume that all pulsars have the same mass and we find that both quantities are conserved with a relative error of $\lesssim 10^{-7}$. The output of the dynamical evolution consists of the position and velocity of each neutron star computed in both galactocentric (GC) and equatorial ICRS (International Celestial Reference System) frames. To transform between different coordinate systems, we employed the method `coordinates` from the Python library `Astropy` (Astropy Collaboration et al., 2013; Astropy Collaboration et al., 2018), where we adopted a galactocentric distance of $r_\odot = 8.3$ kpc and Galactic height of $z_\odot = 0.02$ kpc for the Sun.

5.3   Machine-learning set-up

For our analysis, we will focus on artificial neural networks (ANNs). ANNs are algorithms inspired by the structure of biological brains that can be thought of as nets of interconnected neurons that exchange information from one to another (see Section 2.2). When the network receives an input, it is able to process it to produce an output, like a biological brain responds to external stimulation. In ANNs, neurons are usually organised in a stack of layers. Each neuron in a layer receives input signals (typically real numbers) from the neurons in the previous layer and produces an output signal by applying a non-linear activation function to a linear combination of the input signals according to certain weights and a bias. The output is then passed to the neurons in the following layer, and so on, until the final layer is reached and the output is generated. In our particular case, we will focus on supervised learning, where training the neural network consists of making it produce a specific target output when a particular input is passed through it. This is achieved by (i) labeling the input samples in the training dataset with a label indicating the property that the network has to learn (the so-called ground truth) and (ii) iteratively adjusting the values of weights and biases, also called network parameters, in order to minimise a specific loss function which measures the distance between the network output prediction and the target ground truth (see Section 2.3.4).

Among their numerous applications, ANNs have been employed in regression problems where the network is trained to infer the values of continuous variables for the given input data. This is the kind of problem we are after since we want our

---

[1]`Numba` translates Python functions into optimised machine code at run-time, which allows us to achieve a speed-up by about a factor of 6.
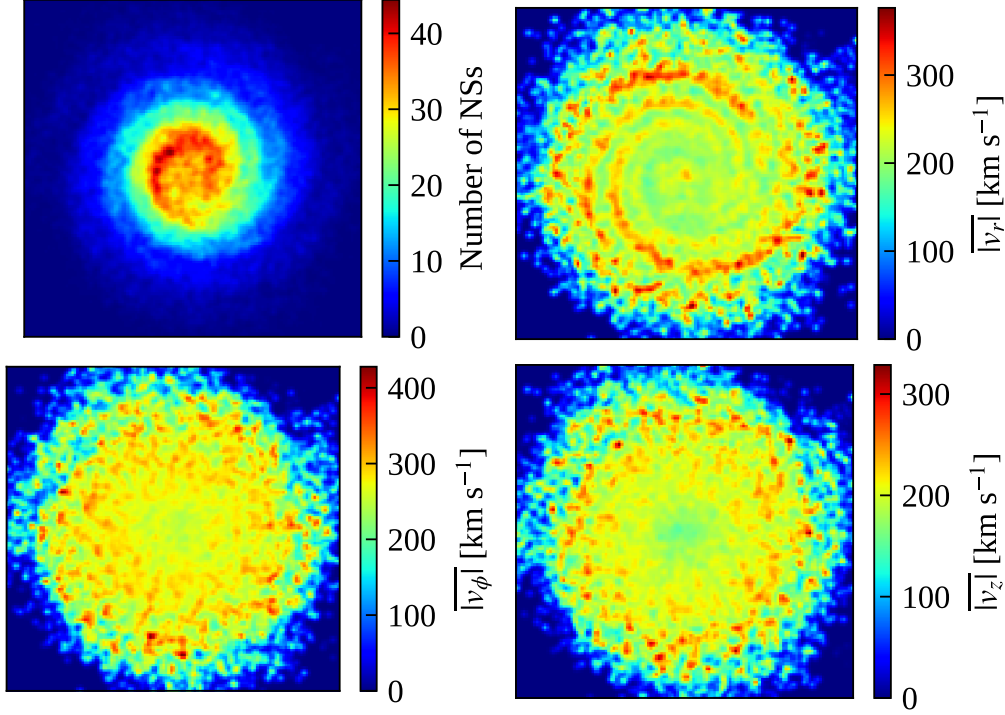
FIGURE 5.1: Examples of 128×128 resolution maps in the galacto-centric xy-plane extending from -20 to $20\,\mathrm{kpc}$ in both x and y direction and showing (in order from left to right) the density of simulated neutron stars, the distribution of average values of the v_r, v_ϕ and v_z velocity components for a population of neutron stars simulated with $h_\mathrm{c} = 0.18\,\mathrm{kpc}$ and $\sigma_\mathrm{k} = 265\,\mathrm{km\,s^{-1}}$. For visualisation purposes, we represent the data using a colourmap to highlight the resulting structures; red regions are characterised by a higher density of stars or higher average magnitude of the velocity components, respectively, while blue areas correspond to lower densities and lower velocity magnitudes. We note that the spiral-arm pattern is still recognisable in the position-density map although high kick velocities tend to blur and disperse the stellar density distribution. In the v_r-velocity map, the inter-arm regions are visible as high-velocity areas, because during the dynamic evolution the space between the spiral arms is progressively filled with high velocity stars that have escaped from their birth places. The other two velocity components exhibit smoother behaviour because the spiral-arm structure is smeared out.
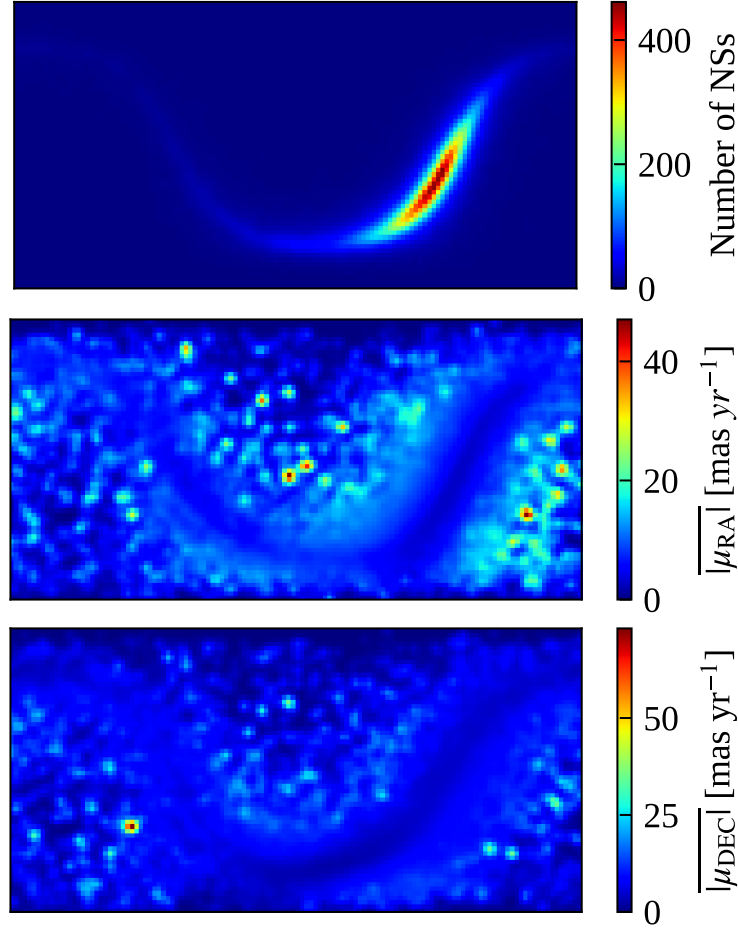
FIGURE 5.2: Examples of 128×64 resolution maps in the equatorial ICRS frame extending from $0°$ to $360°$ in RA and from $-90°$ to $90°$ in DEC and showing (in order from left to right) the density of simulated neutron stars, the distribution of average values for the μ_{RA}, μ_{DEC} proper-motion components for a population of neutron stars simulated with $h_c = 0.18$ kpc and $\sigma_k = 265$ km s^{-1}. For visualisation purposes, we represent the data using a colourmap to highlight the resulting structures; red regions are characterised by a higher density of stars or a higher average magnitude of the velocity components, respectively, while blue areas correspond to lower densities and lower velocity magnitudes. Note that the Galactic silhouette is visible as a *stream* in the position map with an enhanced stellar density close to the Galactic Centre. Due to low-number statistics, the regions outside the Galactic stream in the proper-motion maps are dominated by statistical fluctuations, i.e., the corresponding high-velocity regions are attributed to a small number of high proper-motion stars that have escaped the disk. As a result, the disk itself is dominated by stars with lower proper motion.

network to infer certain parameter values given the evolved neutron star population. In the remainder of this section, we discuss (i) the simulation data we create for our ML experiments (Section 5.3.1), then (ii) focus on the specific network architecture employed (Section 5.3.2), and finally (iii) describe the details of our training process (Section 5.3.3).

5.3.1 Dataset creation and processing

The goal of our ML approach is to predict the parameters that control the dynamical properties of an evolved neutron star population. In particular, we focus on predicting the kick-velocity parameter σ_k and the scale-height parameter h_c, which predominantly affect the distribution of pulsars in the Milky Way. To extract these from an evolved population, and follow a supervised learning approach, we first need to train a neural network on a series of simulated populations (created by exploring the ranges for σ_k and h_c). Following the prescription described in Section 5.2, we perform the following simulation runs:

S1 We generate 10 datasets with an increasing number of samples (specifically 4, 8, 16, 32, 64, 128, 256, 512, 1024 and 20000 simulated populations) by uniformly varying the parameter σ_k of the kick velocity distribution in the range $[1, 700]$ km s^{-1}.[2] We also generate a test dataset with 1000 samples, each one simulated with σ_k randomly drawn from a uniform distribution in the same range of values. For these simulations, we keep the characteristic scale of the z-distribution fixed to its fiducial value $h_c = 0.18$ kpc.

S2 We fix the kick-velocity parameter to its fiducial value $\sigma_k = 265$ km s^{-1} and generate a dataset of 20000 samples of simulated populations by uniformly varying the scale-height parameter h_c in the range $[0.02, 2]$ kpc. We also generate a test dataset with 1000 samples, each one simulated with h_c randomly drawn from a uniform distribution in the same range of values.

S3 We generate 6 datasets, where we uniformly vary the kick-velocity parameter σ_k in the range $[1, 700]$ km s^{-1} *as well as* the characteristic scale of the z-distribution h_c in the range $[0.02, 2]$ kpc. We choose the dataset sizes $16 = 4 \times 4$, $64 = 8 \times 8$, $256 = 16 \times 16$, $1024 = 32 \times 32$, $4096 = 64 \times 64$ and $16384 = 128 \times 128$ given by all the combinations of σ_k and h_c values. As an example: the 16 populations in the first set are obtained by combining each of the 4 values of the σ_k parameter with all 4 values of the h_c parameter. We also generate a test dataset with 1000 samples, each one simulated with both σ_k and h_c randomly drawn from uniform distributions in their respective parameter ranges specified above.

As addressed in detail in Section 5.4, the smaller simulation datasets will be used to explore the network behaviour. The largest datasets containing 20000 and 16384 samples, respectively, and the test datasets with 1000 samples will be used to perform the final training experiments and assess the actual network accuracy in generalisation scenarios.

After the runs have been performed, we transform the output of the simulation into a representation that can be interpreted by a ML pipeline. Since ANNs require the use of structured data, we represent the position and velocity output

---

[2]To generate our simulation data, we partially employ the package Hydra (https://hydra.cc/, Yadan 2019), which allows us to easily sweep entire parameter ranges.

of the simulations in the form of 2D binned density and velocity maps both in the galactocentric and ICRS reference frames. The density maps give information about the density of neutron stars in the Galaxy by providing the number count of stars in each spatial bin. On the other hand, velocity maps contain information about the kinematic properties of the neutron stars by providing the average magnitude of the stellar velocity components in each spatial bin. In the galactocentric maps the Galaxy is represented face on and projected onto the xy-plane of the Cartesian galactocentric frame, extending from -20 kpc to 20 kpc in x and y direction. The ICRS maps instead extend from $0°$ to $360°$ in right ascension (RA) and from $-90°$ to $90°$ in declination (DEC). To each map we apply a smoothing Gaussian filter (with radius 4σ and $\sigma = 1$) in order to add some blurring and avoid sharp features. By doing so, we reduce noisy high-frequency features and thus make the training more stable, presumably resulting in better generalisation capabilities. Therefore, for each simulated population we have:

- 1 density map in the galactocentric frame.

- 3 velocity maps, one for each component of the velocity in cylindrical galactocentric coordinates v_r, v_ϕ and v_z in [km s^{-1}].

- 1 density map in ICRS coordinates.

- 2 proper motion maps, one for each component of the angular proper motion projected on the celestial sphere μ_{RA} and μ_{DEC} in [mas yr^{-1}].

This set of maps for each population is labeled with the corresponding values of the parameters σ_{k} and h_{c} used to simulate that specific population.

We will test the ML performance on three different map resolutions. Thus, we generate each of the above datasets with resolutions 32×32, 128×128 and 512×512 square bins. Note that in the ICRS maps the DEC coordinate axis range is half that of the RA coordinate axis. Hence, these maps have half the bins along the DEC coordinate and their resolution is 32×16, 128×64 and 512×256 square bins. For brevity hereafter we will refer to the three resolutions as 32, 128 and 512 resolutions for both galactocentric and ICRS maps, respectively. An example of the maps with resolution 128 for a simulation with fiducial values $h_{\mathrm{c}} = 0.18$ kpc and $\sigma_{\mathrm{k}} = 265$ km s^{-1} is shown in Figures 5.1 and 5.2.

Before loading the maps into our ML pipeline, they are normalised so that each bin contains a continuous value between 0 and 1. The same applies to the related labels so that their values range continuously between 0 and 1. The aim of normalisation is to speed up the training process and make convergence easier since all inputs will provide signals of similar magnitude to the loss-function minimisation (see Section 2.3.1). This is useful especially for multi-parameter and multi-channel training, that is when we train a network to predict more than one parameter or use channels that have different absolute magnitudes. In these cases, training without normalisation might lead to slower, worse or even no convergence at all. Apart from the blurring and normalisation described above, we do not apply any additional preprocessing steps to our input data.

5.3.2   Network architecture

For the implementation of our ML pipeline we use `PyTorch` (Paszke et al., 2019), an optimised tensor library for deep learning using GPUs and CPUs, that is written in `Python`. The simplest neural network that can be used for this task is a fully
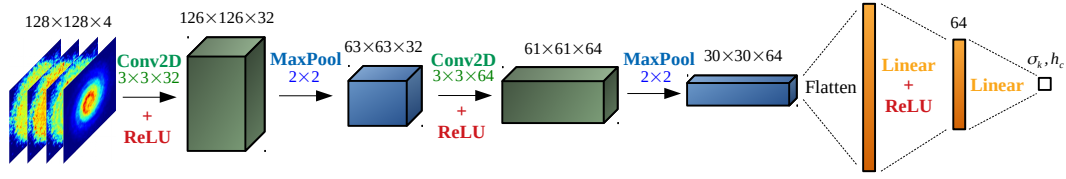
FIGURE 5.3: Schematic representation of our CNN architecture for an input of galactocentric maps with resolution 128×128 and 4 channels (1 density map plus 3 velocity maps). A module formed by two blocks that each contain a convolution layer and max pooling layer is followed by a fully connected linear network with one hidden layer. The final network output is either a single parameter or two parameters depending on the experiment specifics.

connected neural network also referred to as a multi-layer perceptron (MLP) (see Section 2.2.1) with only two layers of neurons, which are referred to as the input and output layer, respectively. As this is the starting point to develop more complicated and advanced architecture models, we first test how this simple configuration behaves. The number of neurons for the input layer is equal to the total number of input features, i.e., $C \times W \times H$. Here, C is the number of input channels (corresponding to the total number of maps used), while W and H are the number of bins in width and height, respectively. The number of neurons in the output layer is equal to the number of regression parameters that we would like to predict, i.e., 1 or 2 in our experiments. For the activation function we use the Rectified Linear Unit (ReLU) defined as $\text{ReLU}(x) = \max(0, x)$ (see Section 2.2.3). To obtain the output values, the ReLU activation function is applied to a linear combination of the input features with weights and a bias.

A more sophisticated model architecture is represented by a convolutional neural network (CNN). CNNs are a particular type of deep neural network that have proven to be very successful in regression and classification tasks when applied to structured and matrix-like 2D inputs (see Section 2.2.2 and Rawat and Wang, 2017, for a review). The basic structure of CNNs consists of convolutional, pooling, and fully connected layers. Convolutional layers are multi-channel filters that slide along the 2D input maps and are able to extract feature maps. The role of the pooling layer is to down-sample the output of a convolutional layer. This inevitably causes a loss of information but in general helps to improve the training efficiency by increasing the size of the receptive field (i.e., the region of the input that produces the feature for each neuron) and reducing the number of trainable parameters. The fully connected layers collect all the output features from the convolution layers into a 1D input and return the final output prediction.

The detailed structure of the CNN we built for our case study can be found in Table 5.3. A schematic representation of its structure for a 4-channel input with galactocentric maps is also shown in Figure 5.3 as an example. It consists of the following layers:

- A 2D convolution layer with kernel size 3×3, C input channels, 32 output channels, stride 1 and no padding.

- A 2D Max pooling layer of size 2×2 with stride 2 and no padding.

- A 2D convolution filter with kernel size 3×3, 32 input channels, 64 output channels, stride 1 and no padding.

TABLE 5.3: CNN architecture. The total number of input and output features is reported. C is the number of used channels, while W and H represent the number of bins (i.e. the resolution) in width and height of the density and velocity maps, respectively. Input and output feature numbers have been rounded down to the lower integer.

Layer	Input	Output
Conv2d + ReLU	$C \times W \times H$	$32 \times (W - 2) \times (H - 2)$
MaxPool2d	$32 \times (W - 2) \times (H - 2)$	$32 \times \left(\frac{W}{2} - 1\right) \times \left(\frac{H}{2} - 1\right)$
Conv2d + ReLU	$32 \times \left(\frac{W}{2} - 1\right) \times \left(\frac{H}{2} - 1\right)$	$64 \times \left(\frac{W}{2} - 3\right) \times \left(\frac{H}{2} - 3\right)$
MaxPool2d	$64 \times \left(\frac{W}{2} - 3\right) \times \left(\frac{H}{2} - 3\right)$	$64 \times \left(\frac{W}{4} - \frac{3}{2}\right) \times \left(\frac{H}{4} - \frac{3}{2}\right)$
Linear + ReLU	$64 \times \left(\frac{W}{4} - \frac{3}{2}\right) \times \left(\frac{H}{4} - \frac{3}{2}\right)$	64
Linear	64	1(2)

- A 2D Max pooling layer of size 2×2 with stride 2 and no padding.

- A fully connected linear layer with flattened input from the convolutional modules' output and 64 output neurons.

- A fully connected linear layer with 64 input neurons and 1 or 2 output neurons (depending on the number of parameters we would like to predict).

For the convolutional and pooling layers the stride parameter regulates the amount of displacement in bins with which the filter moves over the map at each step. Padding adds one or more bins at the border of the 2D maps, so that the filters can move and cover the whole map without leaving any bins out. We use a padding of 0 because the borders of the maps do not contain relevant information.

The choice of this architecture was found by trial and error experiments where we started from a very simple structure and progressively increased the complexity, adding more and more layers to acquire the desired accuracy in predicting the input parameters.

5.3.3 Training process

For the training of the network, we use the root-mean-square error (RMSE) both for the loss function and validation metric, i.e., to compute the distance between the network predictions and the ground truths of the h_c and σ_k parameters. In general, validation occurs at the same time as training and consists of testing the network over a dataset different from the training set. This is needed to asses the ability of the network to generalise what it is learning to an unknown dataset (see Section 2.3). The minimisation of the loss function occurs through gradient descent and backpropagation (Kelley, 1960; Ruder, 2017), i.e., computation of the loss-function gradients with respect to all network parameters (weights and biases). These gradients taken with a negative sign indicate the directions towards which the network parameters should be updated so that the loss is reduced, and hence the network predictions move closer to the true, expected labels. In this regard, a crucial aspect to ensure the best performance of a neural network is to properly initialise the weights and biases (see Section 2.3.2). For this purpose we use the Kaiming initialisation method (He et al., 2015) in order to avoid exploding or vanishing gradients during the training.

The training process itself is regulated by several hyperparameters. The first one is the learning rate, which is a coefficient for the weight updates. In general, a larger learning rate results in updates of larger magnitude, which could in turn lead to faster convergence, but might also reduce the stability of the training process and thus increase the risk of overshooting the minima of the loss landscape. A second hyperparameter is the batch size, which defines the number of samples that are packed together and passed through the network before an optimisation step is performed. In general, for bigger training datasets, a larger batch size helps to increase the efficiency and stability of the training process, since the gradient-descent steps are averaged over many samples and noise is reduced (see Section 2.3.3). For the gradient-descent optimiser, we use the Adaptive Moment Estimation (Adam) (Kingma and Ba, 2014). As its name suggests, Adam adds an adaptive momentum term to the gradient descent to automatically modify the learning rate and accelerate convergence (see Section 2.3.4). When using the Adam optimiser, the chosen initial value of the learning rate represents only an upper limit.

We fix the maximum number of learning epochs to 1024. Every epoch the network performs a series of optimisation steps by going through the whole training dataset once. Then epoch-averaged loss and validation metric values are computed. If the validation metric value has improved with respect to the previous epoch the current status of the optimised network is saved. We set an early stop of 128 epochs, so that if the validation metric does not improve over this epoch span, the training process automatically stops and the weights of the best epoch are stored. This prevents the network from overfitting the training samples, which would reduce its ability to generalise over unknown data.

5.4 Experiments

For our machine-learning experiments we use the hardware and software specified in Appendix A.1. As a first step, we perform several tests to analyze which configuration of the input feature maps provides the best training experience and which proposed neural network architecture, either the MLP or the CNN, behaves better. Regarding the input configuration, we would like to understand (i) what the best type of maps is (galactocentric vs equatorial ICRS); (ii) how many input channels are needed to obtain good results (do density maps provide enough information alone or does the addition of velocity maps improve the results?); (iii) which resolution of the input maps provides the best result. To do so, we first compare the behaviour of the MLP and the CNN when trained to predict a single parameter. We explore different types of input signals by varying the resolution of the density and velocity maps as well as the number of input channels. Once we find the optimal configuration for the input maps and the best performing network, we proceed to test its generalisation power. A similar strategy is then followed for the two-parameter prediction.

5.4.1 Single-parameter predictions

Data-representation and architecture comparison

We focus on predicting the parameter σ_k of the Maxwell kick velocity function and employ the density and velocity maps generated from simulation run S1. We keep aside the dataset with 20000 samples as it will be used to asses the generalisation power of the best performing network (see Section 2.3). Thus, we have training sets

FIGURE 5.4: MLP best training RMSE values for the training process on the single parameter σ_k of the Maxwell kick-velocity distribution, as a function of the training dataset size and the resolution (*red*, *blue*, and *orange* curves for 32, 128 and 512, respectively) using the four different input configurations T1 (GC position), T2 (GC position + velocity), T3 (ICRS position) and T4 (ICRS position + velocity).

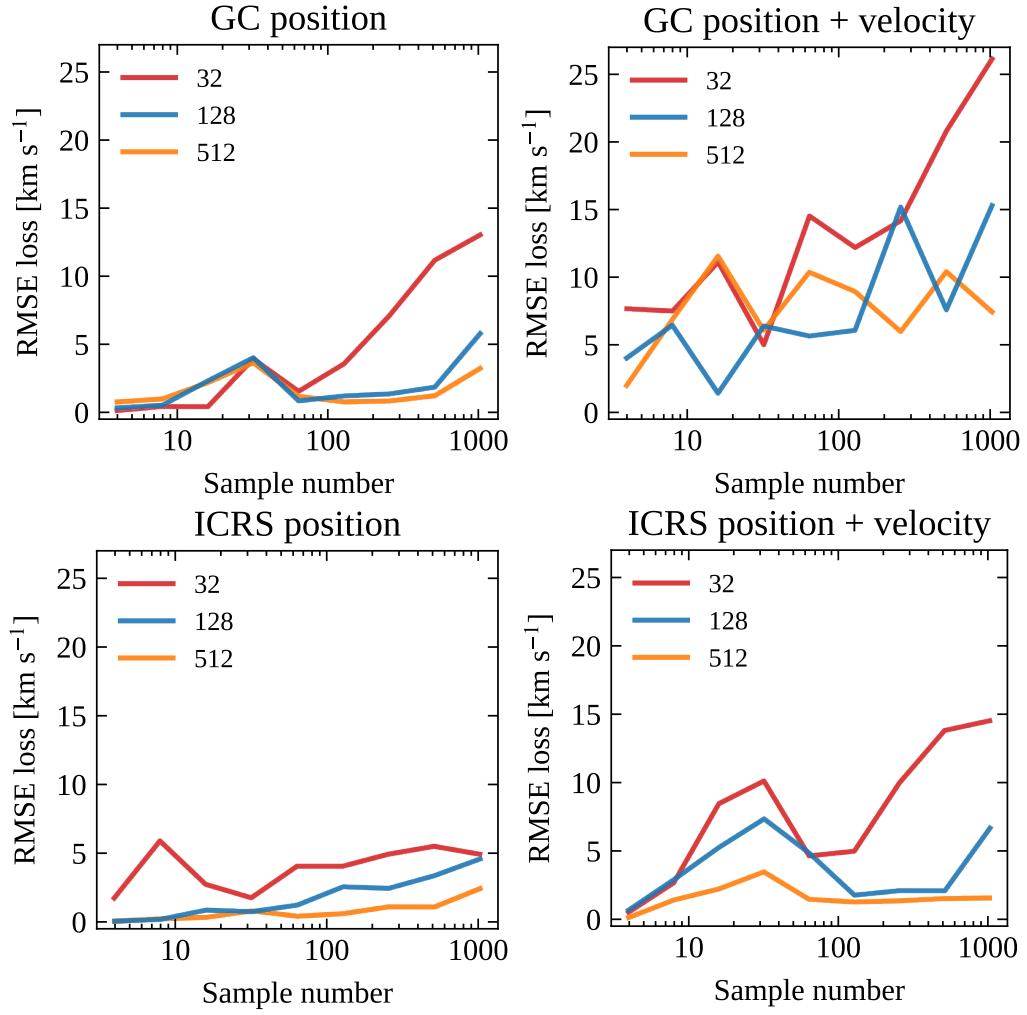FIGURE 5.5: CNN best training RMSE values for the training process on the single parameter σ_k of the Maxwell kick-velocity distribution, as a function of the training dataset size and the resolution (*red*, *blue*, and *orange* curves for 32, 128 and 512, respectively) using the four different input configurations T1 (GC position), T2 (GC position + velocity), T3 (ICRS position) and T4 (ICRS position + velocity).
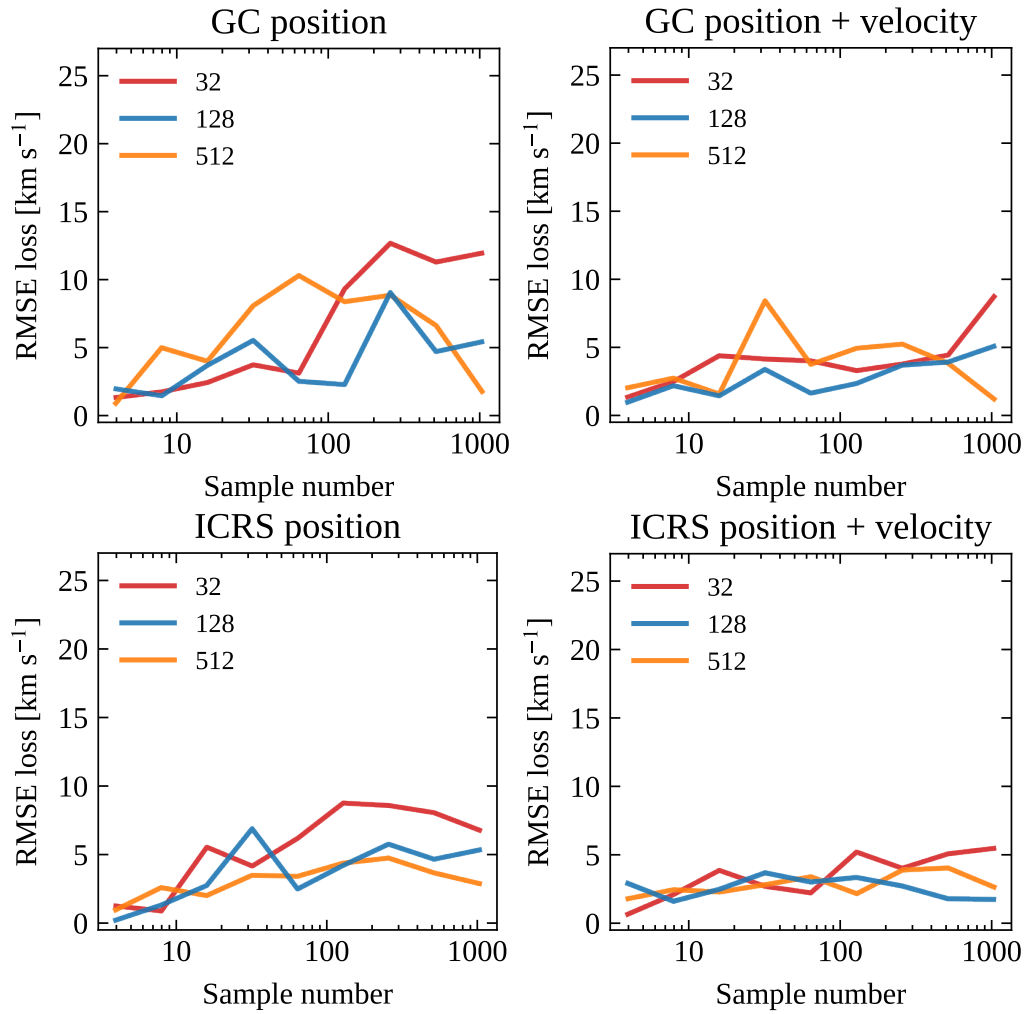
with increasing number of samples (from 4 to 1024) and increasing map resolution (32, 128, 512). For each of these training sets, we compare the performance on four different kinds of input information:

T1 Galactocentric position information: 1 density map in galactocentric coordinates (1 channel).

T2 Galactocentric position and velocity information: 1 density map plus 3 velocity maps in galactocentric coordinates (4 channels).

T3 ICRS position information: 1 density map in equatorial ICRS coordinates (1 channel).

T4 ICRS position and velocity information: 1 density map plus 2 proper motion maps in equatorial ICRS coordinates (3 channels).

We fix the network architectures and the training set-up as described in Section 5.3.2 and Section 5.3.3. For these initial experiments, we do not incorporate validation but only focus on the training results for different types of datasets, as we are not yet interested in evaluating the generalisation power of the networks. To assess the convergence of the training runs, we set a threshold for the σ_k RMSE training loss to $10 \, \mathrm{km \, s^{-1}}$. If the final RMSE value is higher than this threshold the training is repeated up to a maximum of 8 times. If convergence is not reached after 8 trials we take the trained model with the lowest final RMSE. For each training experiment, we also monitor the computational time needed to perform a single optimisation step on a single data batch (see Appendix A for more details).

Initially, we perform a search for the starting value of the learning rate that provides the best results. In particular, for the MLP we find that to ensure a decaying RMSE value during training, the initial learning rate needs to be decreased as the input-map resolution increases. Therefore, after several tests, we set the initial learning rate to 10^{-4}, 10^{-5} and 10^{-6}, respectively, for the 32, 128 and 512 resolution maps. The CNN instead is more flexible and stable and all three initial learning rates are suitable for every resolution. In this case, we set it to 10^{-4} to obtain training convergence in the smallest number of epochs.

Next, we tune the batch size to make the learning process more stable and efficient as the number of training samples increases. In general, we keep the batch size to 1 for the dataset sizes 4, 8, 16, 32 and progressively increase it to 4, 8, 16, 32, 64 as the dataset size increases to 64, 128, 256, 512, 1024, respectively. Only when testing the performance of the MLP on the T2 input configuration, we need to further fine tune the batch size in order to reach acceptable values of the RMSE. In all other cases, the batch sizes mentioned above work well.

The results of our experiments using the MLP and the CNN on the training datasets from S1 with configuration T1, T2, T3 and T4 are shown in Figures 5.4 and 5.5, respectively (see Appendix A for the timing results), highlighting the best converged RMSE values as a function of the training dataset size and the resolution.

First of all, we note that for both model architectures, ICRS maps allow us to obtain slightly better results in terms of the best RMSE values. An explanation for this could be that only the ICRS maps contain 3D information of the Galaxy, i.e., they encode the stellar height distribution with respect to the Galactic disk which correlates with the kick-velocity magnitude. In fact, the higher the kick velocity of newborn neutron stars the more spread out their distribution in galactic height z at the end of the dynamical evolution. On the other hand, galactocentric maps show the Galaxy represented face on and the information on the stars' height z with
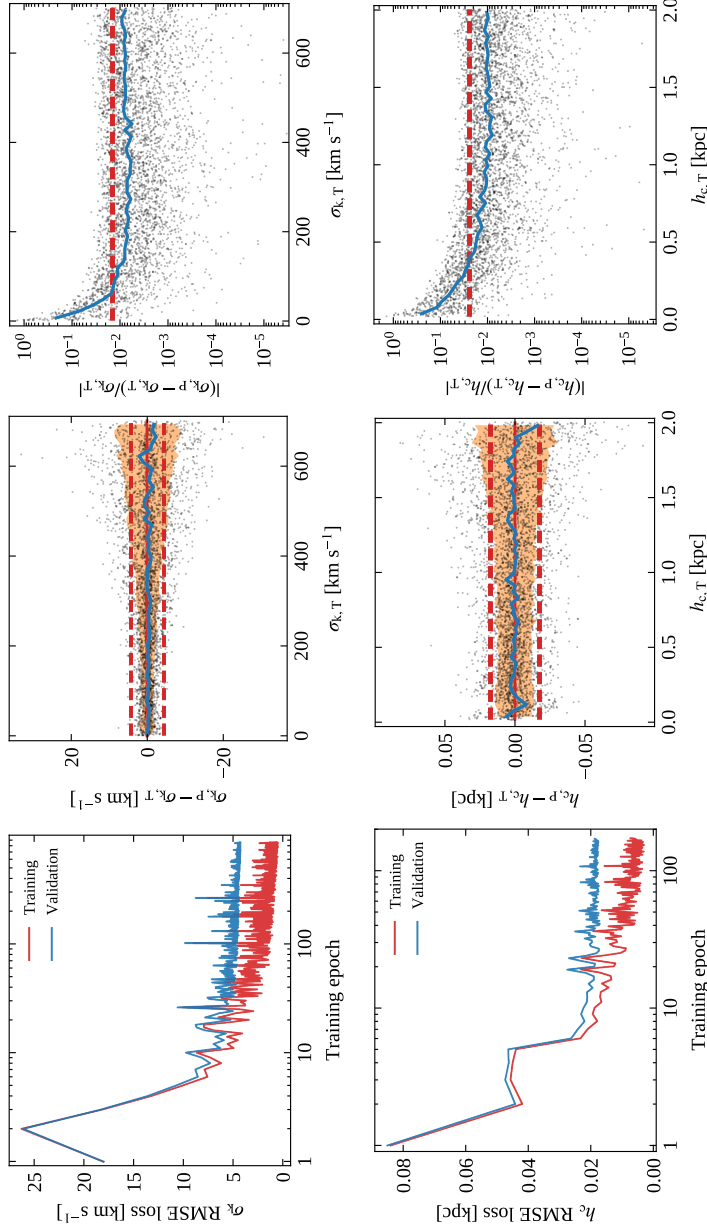
FIGURE 5.6: Results of the CNN single-parameter prediction for the kick-velocity parameter, σ_k, and scale height, h_c, for the corresponding validation sets. *Top (bottom) left panel*: evolution of the RMSE training (*red*) and validation (*blue*) losses as a function of the training/validation epoch for the σ_k (h_c) parameters. *Top (bottom) central panel*: residuals of the prediction as a function of the target σ_k (h_c) value; the subscripts P and T refer to the predicted and target values, respectively. The red dashed lines delimit the 68% uncertainty region corresponding to a RMSE = $4.4 \, \mathrm{km \, s^{-1}}$ ($0.017 \, \mathrm{kpc}$) computed over the whole range $[1,700] \, \mathrm{km \, s^{-1}}$ ($[0.02, 2] \, \mathrm{kpc}$). The *orange* region delimits the 68% uncertainty region computed as a running RMSE for which we have divided the full σ_k (h_c) range into 50 bins of equal size. The *blue* line shows the trend of the average residuals which are well centred around the value 0. *Top (bottom) right panel*: relative error of the prediction as a function of the target σ_k (h_c) value. The red dashed line corresponds to a MRE = 0.014 (0.024) computed over the whole range $[1,700] \, \mathrm{km \, s^{-1}}$ ($[0.02, 2] \, \mathrm{kpc}$). The *blue* line shows the trend of the running MRE computed over 50 bins of equal size into which we have divided the full σ_k (h_c) range.

respect to the Galactic plane is hidden. Therefore, it is likely easier for the networks to distinguish populations with different σ_k by processing the ICRS maps.

We also note that the addition of the velocity or proper motion information has opposite effects on the two model architectures. In particular, on one side it worsens the MLP performance, as the best RMSE almost doubles. On the other side, it helps the CNN in reducing the overall RMSE. We interpret this as an indication that, as the complexity of the input data increases, a deeper (with more layers) and more sophisticated model architecture like the CNN is more suitable to process the input data and extract meaningful features to perform the regression task. However, the improvement is not dramatic, which could suggest that the density maps already provide enough information to distinguish, with sufficient precision, populations simulated with different σ_k values.

Concerning the resolution, we observe that higher resolutions allow us to reach slightly lower RMSE values with both the MLP and the CNN but at the expense of longer computation time (see Figures A.1 and A.2 in Appendix A). We therefore consider the small differences between the best RMSE values obtained with 128 and 512 resolutions not sufficient to justify the choice of the higher resolution. Hence, the use of the 128 resolution appears to be a good compromise to ensure good accuracy and reasonably fast training.

In light of these results for the single-parameter estimation, we conclude that the optimal representation to be used for training is composed of the ICRS density plus proper motion maps with 128 resolution. Moreover, as the CNN obtains the best results and appears more stable and flexible when compared to the simple MLP (especially for the multi-channel input features), we employ our CNN architecture in the following experiments, which should also be less prone to overfitting.

Generalisation Results

As the next step, we separately train the CNN to predict the σ_k and h_c parameters by using the two big datasets with 20000 simulations each (see S1 and S2). As input features, we use the 3-channel ICRS representation with one density map and two proper motion maps with 128 resolution that ensure the best results as suggested by our earlier experiments. We randomly split both datasets into training and validation subsets with a relative percentage of 80/20%, respectively. Therefore, training is performed over 16000 simulations, randomly sampled from the entire datasets, while validation is performed over the remaining 4000 simulations. We adopt an initial learning rate of 10^{-4}, a batch size of 64, set the total number of learning epochs to 1024 and an early stop at 128 epochs to avoid overfitting. We set the convergence threshold for σ_k to $10\,\mathrm{km\,s^{-1}}$ and for h_c to $0.5\,\mathrm{kpc}$. The evolution of the training and validation losses is shown in the left panels of Figure 5.6.

The predictions of the trained network on the validation set for σ_k and h_c are summarised in the central panels of Figure 5.6 and Table 5.4. In the first case, the network is able to predict the value of the kick-velocity parameter σ_k for the simulations in the validation dataset with a RMSE uncertainty of $4.4\,\mathrm{km\,s^{-1}}$, computed over the whole range $[1, 700]\,\mathrm{km\,s^{-1}}$. This is indicated by the red dashed lines in the residuals plot (see top central panel of Figure 5.6), which delimit the 68% uncertainty region. In the second case, the network is able to predict the value of the scale-height parameter h_c for the simulations in the validation dataset with a RMSE uncertainty of $0.017\,\mathrm{kpc}$, computed over the whole range $[0.02, 2]\,\mathrm{kpc}$ (see bottom central panel of Figure 5.6). Note that in both experiments, the residuals spread out as the target values increase. We visualise this by computing a running RMSE with

increasing target values. As is shown by the orange regions in the residuals plots, the running RMSE increases from 1.4 to 7.1 km s^{-1} for σ_k and from 0.013 to 0.028 kpc for h_c, respectively. However, the average residuals are consistent with 0 over the entire target value ranges as marked by the blue line in the residual plots, showing no anomalous trends in the predicted values.

In the right panels of Figure 5.6, we also show the relative error to highlight the precision with which the network is able to predict the σ_k and h_c parameters for a given target value. We observe that the precision of the predictions improves with increasing σ_k and h_c and eventually stabilises to a relative error of around 0.01. This is highlighted by the blue lines, which show the trend of the mean relative error (MRE) as a function of the target parameters. The red dashed lines correspond instead to the MRE computed on the whole parameter ranges and are equal to 0.014 and 0.024 for σ_k and h_c, respectively. The fact that the precision of the predictions decreases at the lower end of the target ranges suggests that (for our chosen network set-up) the input maps become harder to distinguish as the neutron stars' initial kick velocities and their galactic birth heights decrease in magnitude.

We then evaluate the generalisation capability of the two trained networks on the corresponding test-sets with 1000 samples each. We find RMSEs of 4.8 km s^{-1} and 0.019 kpc and MREs of 0.018 and 0.029 for σ_k and h_c, respectively. To further assess the confidence intervals of both estimators, we evaluate the RMSE and the MRE of the parameter values predicted by the two networks over 1000 bootstrapped sets of the related test sets. We find that the RMSE variation is around 3%, while the MRE variation is around 11% for both predicted parameters. These results indicate that the trained networks are able to generalise well over unseen datasets.

5.4.2 Two-parameter predictions

Data-representation comparison

To see how the CNN behaves when two parameters, i.e., the kick-velocity parameter σ_k and the characteristic scale height h_c, are inferred simultaneously, we use the dataset of maps generated from simulation run S3 (see Section 5.3.1). In this case, we have training sets with increasing number of samples, 16 = 4x4, 64 = 8x8, 256 = 16x16, 1024 = 32x32, 4096 = 64x64. We leave aside the largest dataset with 16384 = 128x128 simulations for our final generalisation experiment. Given the results of the single-parameter training experiments, we choose the 128 resolution maps and compare the CNN's performance on the four kinds of input information T1, T2, T3 and T4 as we did for the single-parameter case (see Section 5.4.1). We fix the CNN architecture and the training set-up as described in Sections 5.3.2 and 5.3.3, respectively. As for the single-parameter comparison, we only focus on the training behaviour for this initial comparison. We keep track of the RMSE training losses for both parameters separately, but training proceeds by minimising the total loss computed on both parameters. We set the convergence threshold for σ_k to 10 km s^{-1} and for h_c to 0.5 kpc; if convergence is not reached for both parameters after 8 training trials we quote the experiment with the best performance. As before, we also keep track of the computational time for a single optimisation step (see Appendix A). The initial learning rate is set to 10^{-4}, while the batch size is changed according to the dataset size. In particular, we use a batch size of 1, 4, 16, 64, 128 for the dataset sizes 16, 64, 256, 1024 and 4096, respectively.
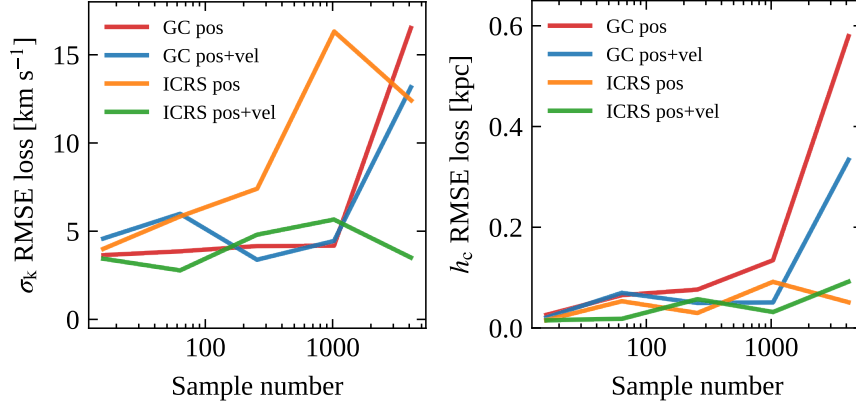
FIGURE 5.7: CNN best training RMSE values for the training process on the two parameters σ_k of the Maxwell kick-velocity distribution (*left panel*) and characteristic scale height h_c (*right panel*) as a function of the training dataset size for the four different input configurations T1 (GC position), T2 (GC position + velocity), T3 (ICRS position) and T4 (ICRS position+ velocity).

The results of our experiments using the CNN for the two-parameter prediction on the training datasets from S3, with configuration T1, T2, T3 and T4 are summarised in Figure 5.7 (see Appendix A for the timing results). As in the single-parameter case, we find that the best results are provided by the ICRS maps, which allow us to reach the lowest training RMSE losses for both parameters. In particular, for the ICRS 3-channel input, the CNN is able to reach a training RMSE loss $\lesssim 5\,\mathrm{km\,s^{-1}}$ for the σ_k parameter, comparable with the single-parameter case. For h_c, the CNN reaches a training RMSE $\lesssim 0.1\,\mathrm{kpc}$. However, we note a drop in accuracy for the σ_k parameter when only the ICRS density maps are used. As already mentioned in Section 5.4.1, information on the stars' z-coordinate is encoded in the ICRS maps. As we simultaneously vary the initial kick velocities and the galactic heights of the pulsars' birth places, the degeneracy between the effects of these two parameters becomes relevant. This makes a distinction of the impact of one parameter over the other more difficult for the network, when only ICRS density maps are provided (see also Section 5.4.2). Adding two extra channels that contain information about the stars' proper motion thus helps to improve the accuracy on the kick-velocity parameter estimation. The results of these initial explorations are promising and indicate that our simple CNN architecture has good predictive power for both parameters, if provided with all three input channels in the ICRS representation.

Generalisation Results

As for the single-parameter analysis, we assess the actual performance of the network when simultaneously predicting σ_k and h_c by training the CNN on the biggest dataset with $16384 = 128 \times 128$ simulations (see run S3). As discussed above, we use the 3-channel ICRS representation with one density map and two proper motion maps with 128 resolution as input features. We split the entire dataset into training and validation subsets with a relative percentage of 80/20%, respectively, leading to 13107 samples in the training and 3277 samples in the validation dataset and use the same configuration as in Section 5.4.1. The evolution of the individual training and validation losses is shown in the left panels of Figure 5.8.
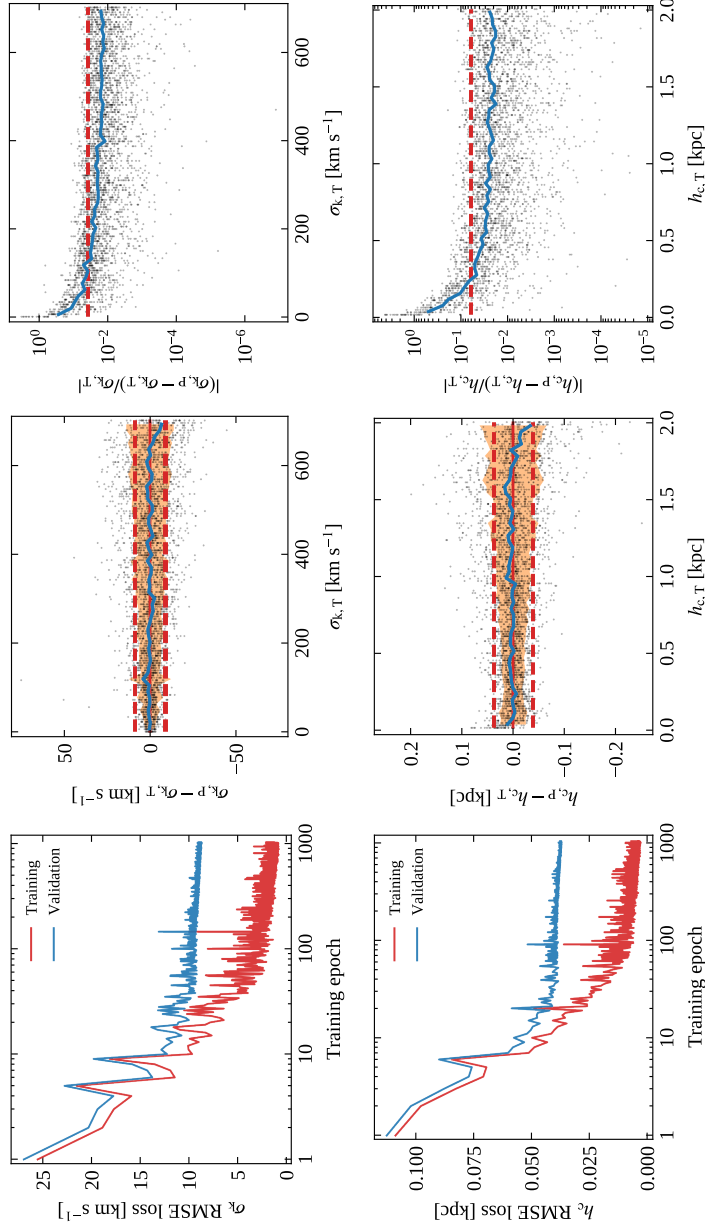
FIGURE 5.8: Results of the CNN two-parameter prediction for the kick-velocity parameter, σ_k, and scale height, h_c, for the corresponding validation sets. *Top (bottom) left panel*: evolution of the RMSE training (*red*) and validation (*blue*) losses as a function of the training/validation epoch for the σ_k (h_c) parameters. *Top (bottom) central panel*: residuals of the prediction as a function of the target σ_k (h_c) value; the subscripts P and T refer to the predicted and target values, respectively. The red dashed lines delimit the 68% uncertainty region corresponding to a RMSE = 8.8 km s^{-1} (0.038 kpc) computed over the whole range [1,700] km s^{-1} ([0.02,2] kpc). The *orange* region delimits the 68% uncertainty region computed as a running RMSE for which we have divided the full σ_k (h_c) range into 50 bins of equal size. The *blue* line shows the trend of the average residuals which are well centred around the value 0. *Top (bottom) right panel*: relative error of the prediction as a function of the target σ_k (h_c) value. The red dashed line corresponds to a MRE = 0.039 (0.061) computed over the whole range [1,700] km s^{-1} ([0.02,2] kpc). The *blue* line shows the trend of the running MRE computed over 50 bins of equal size into which we have divided the full σ_k (h_c) range.

TABLE 5.4: Summary of the CNN generalisation results on the validation and test (in parenthesis) datasets for the single-parameter and two-parameter cases.

Parameter	1-par. generalisation		2-par. generalisation	
	RMSE	MRE	RMSE	MRE
σ_k	4.4 (4.8) km s^{-1}	0.014 (0.018)	8.8 (9.1) km s^{-1}	0.039 (0.033)
h_c	0.017 (0.019) kpc	0.024 (0.029)	0.038 (0.041) kpc	0.061 (0.057)

The results for the trained network's prediction on the validation set for both parameters are shown in the central panels in Figure 5.8 and summarised in Table 5.4. The network is able to predict σ_k and h_c with an average RMSE uncertainty of 8.8 km s^{-1} and 0.038 kpc, respectively, which is approximately doubled compared to the single-parameter experiment. These RMSE values are computed over the full target ranges and represented by the red dashed lines in the residuals plots (see central panels of Figure 5.8). As in the single-parameter case, we observe that the RMSE uncertainties increase with increasing target parameters as indicated by the orange regions in the residuals plots. The relative errors represented in the right panels of Figure 5.8 show the same decreasing trend with the target value as for the single-parameter predictions, albeit with larger relative errors. When computing the MRE over the whole range of the two target parameters, we obtain 0.039 and 0.061 for σ_k and h_c, respectively. These values are highlighted by the red dashed lines in the right panels of Figure 5.8.

We then evaluate the generalisation capability of the trained network on the test set with 1000 samples. We find RMSEs of 9.1 km s^{-1} and 0.041 kpc and MREs of 0.041 and 0.057 for σ_k and h_c, respectively. As before, we evaluate the confidence intervals of the two estimators over 1000 bootstrapped sets of the test set and find that the RMSE and MRE variations are around 3% and 8%, respectively, for both predicted parameters. This indicates that also in the two-parameter prediction the trained network is stable and guarantees a good level of generalisation power.

However, we find that the CNN trained to simultaneously predict σ_k and h_c is not able to reach the same level of accuracy as in the experiments where a single parameter was predicted at a time. This could be due to three distinct causes: (i) either our neural network is not sophisticated enough to discern between the effects of both parameters on the simulation outcomes represented in the ICRS maps, (ii) our choice of ICRS maps as input does not provide sufficient information for the network to distinguish between both parameters, or (iii) this is a physical (real) degeneracy, and there is a limit to what we can measure. To investigate this issue, we train the CNN to predict *only* the parameter h_c using the same two-parameter dataset used above where both σ_k and h_c are varied. After predicting on the validation dataset, we obtain a RMSE accuracy of 0.038 kpc which is equal to the result obtained above for the two-parameter prediction. This suggests that the network complexity is suitable to predict either one or two parameters simultaneously. Limitations in performance are therefore either due to an inadequate input representation or a physical degeneracy that imposes a natural accuracy threshold. While we cannot distinguish these two with our current simulation and ML pipeline, we can illustrate the underlying problem in the following way: Figure 5.9 shows the residuals of the scale-height parameter h_c versus the residuals of the kick-velocity parameter σ_k for the predictions
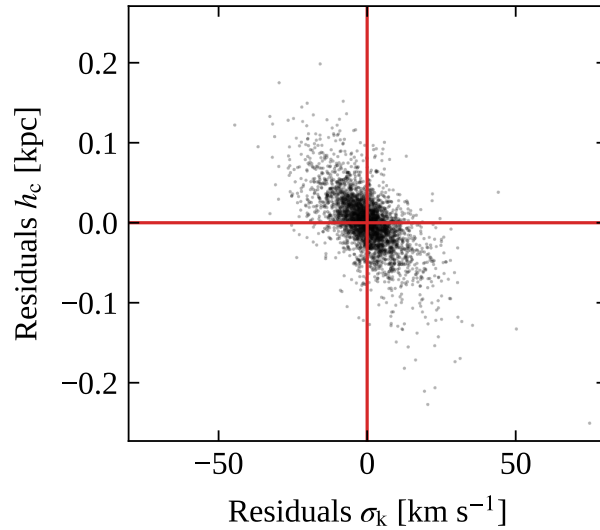
FIGURE 5.9: Scatter plot of residuals of the predicted scale-height parameter h_c versus residuals of the predicted kick-velocity parameter σ_k for the validation dataset of our two-parameter generalisation experiment. An anticorrelation can be observed.

over the validation set. The overall negative slope indicates that the network tends to overpredict h_c in those simulations where σ_k is underestimated and vice versa; i.e., large (small) h_c values have the same overall effect on the phenomenology of the pulsar population as large (small) σ_k values and the network struggles to distinguish these cases. This highlights the degeneracy between the two parameters already discussed above, which might be broken if the data itself were represented in a different way or additional input information about each neutron star (beyond position and velocity) were provided.

5.5 Discussion

In this chapter, we have studied the potential of an artificial neural network to estimate with high accuracy the dynamical characteristics of a mock population of isolated pulsars. Implementing a simplified population-synthesis framework, we focused on the pulsar natal kick-velocity distribution and the distribution of birth distances from the Galactic plane. Taking into account the Galaxy's gravitational potential and evolving the pulsar motions in time, we generate a series of simulations that are used to train and validate a suitably structured convolutional neural network.

The generalised results presented in the previous sections are obtained in a very idealised and simplified scenario, implying that caution is required when the uncertainties for the prediction of the kick-velocity dispersion, σ_k, and birth scale height, h_c, are taken at face value and conclusions for the real pulsar population are drawn. In particular, our simulations assume that the distribution of neutron-star progenitors in Galactic height is represented by the exponential thin-disk model, and that the kick-velocity magnitudes follow a Maxwellian distribution. While the choice of an exponentially thin disk is commonly adopted (Wainscoat et al., 1992; Polido, Jablonski, and Lépine, 2013; Li et al., 2019) and can be justified theoretically as the outcome of a self-gravitating isothermal disk (Spitzer, 1942), the choice of a Maxwellian model for the kick-velocity distribution is difficult to motivate from a
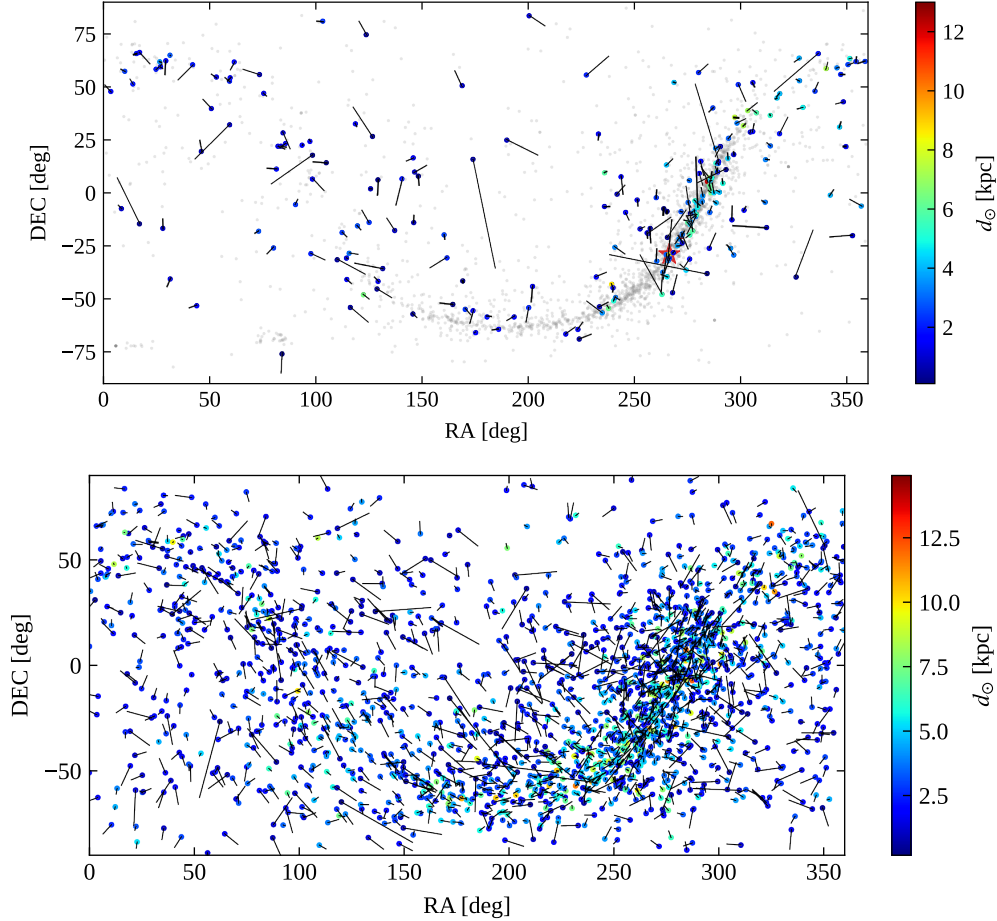
FIGURE 5.10: Proper motion trajectories in the ICRS reference frame of the selected observed 216 neutron stars (*upper panel*) and 2000 mock neutron stars simulated using the fiducial values for the kick velocity and scale height ($\sigma_\text{k} = 265\,\text{km}\,\text{s}^{-1}$ and $h_\text{c} = 0.18\,\text{kpc}$) and sampled using the wight function $f(d_\odot)$ (*lower panel*). The current locations of the neutron stars are indicated by the coloured circles, whereas the tracks indicate their motion for the past $0.5\,\text{Myr}$, assuming no radial velocity and neglecting the effects of the Galactic potential. The colour encodes the heliocentric distance, $d_\odot$, of the neutron stars. The corresponding data is provided in Table B.1. In the background, we show in grey all non-binary pulsars in the ATNF catalogue (those in the Small and Large Magellanic Clouds as well as those in globular clusters are included). The red star highlights the position of the Galactic Centre.

theoretical standpoint. The Maxwellian model has found empirical support as it has been shown to well describe the proper motions of observed pulsars (Hobbs et al., 2005). It is for this reason, and its rather simple mathematical form, that the Maxwell kick-velocity distribution is often adopted in population synthesis studies of compact stars (Sartore et al., 2010; Cieślar, Bulik, and Osłowski, 2020). However, the real functional form of the kick-velocity distribution is still unknown and debated. Several authors have studied the kick-velocity problem and concluded that other models explain observed proper motions of neutron stars equally well. For example, by using maximum-likelihood methods Arzoumanian, Chernoff, and Cordes (2002) found that a bimodal Gaussian distribution with a low-velocity and a high-velocity component is the preferred model to describe the observed proper motion of a sample of 79 radio pulsars. Faucher-Giguère and Kaspi (2006) studied the velocity component along the Galactic longitude for a sample of 34 pulsars observed with interferometric techniques (Brisken et al., 2002; Brisken et al., 2003). After testing a two-component Gaussian model as well as a variety of single-component models, they opted for a single-component description with an exponential shape, although a two-component model could not be ruled out due to the poor statistics of their sample. More recently, Verbunt, Igoshev, and Cator (2017) and Igoshev (2020) analyzed a sample of isolated young pulsars and found that a two-component Maxwellian model explained the observed sample best.

In general, the presence of a low-velocity and a high-velocity component could indicate different progenitor properties as well as birth scenarios for the pulsar population. Numerical simulations of supernova explosions have, for example, suggested that neutron stars with lower kick velocities could be generated in the core-collapse supernovae of progenitors with small iron cores or in electron-capture supernovae (Podsiadlowski et al., 2004; Mandel and Müller, 2020). While also possible scenarios for isolated systems (Janka, 2017), such conditions might generally affect those neutron stars born in binaries (Giacobbo and Mapelli, 2020), where mass-loss episodes could strip their progenitors off their hydrogen envelopes. This might favour the formation of small iron cores or accretion-induced electron-capture supernovae, resulting in weaker natal kicks (Schwab, Podsiadlowski, and Rappaport, 2010; Tauris et al., 2013). Only for the strongest kicks can the binaries be disrupted by the supernova and both companions expelled; otherwise the two stars remain gravitationally bound. Such effects are neglected in our simplified model but could in principle generate an imprint on the observed neutron-star population.

Up to this point, we have not considered any kind of selection effects or observational biases and effectively assumed that all the neutron stars in our simulation are detectable. While this provides direct insight into how various initial conditions affect the evolved population of neutron stars, a direct comparison with observational data in principle requires a careful treatment of biases. For example, due to beaming effects not all Galactic radio pulsars are visible from Earth (Tauris and Manchester, 1998; Melrose, 2017, see also Section 1.9), while survey sensitivity thresholds and instrumental limitations might hamper the detection of faint or far away sources (Manchester et al., 2001; Johnston et al., 2008; Stovall et al., 2014; Coenen et al., 2014; Good et al., 2020, see also Section 1.12). Additionally, timing noise can significantly limit the sensitivity and precision in the detection of pulsar proper motions via timing analysis techniques (Hobbs et al., 2004; Lentati et al., 2016; Parthasarathy et al., 2019). With the aim of obtaining a rough idea on how selection effects and biases
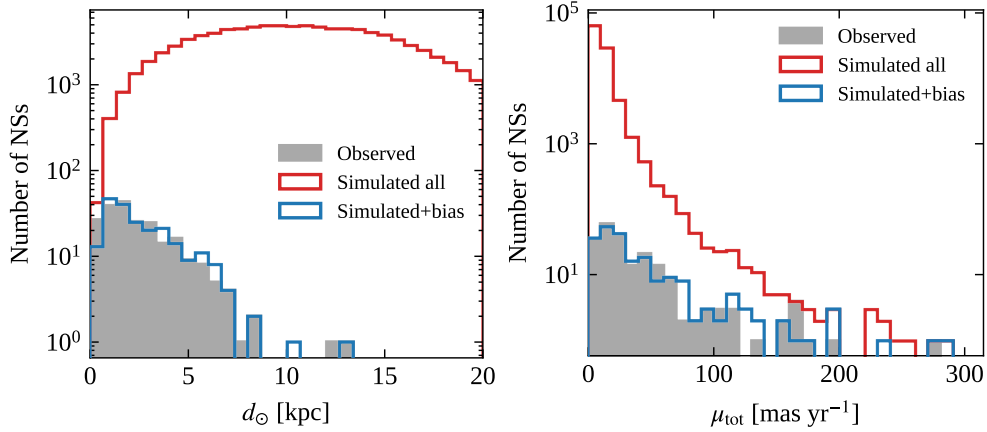
FIGURE 5.11: Distribution of heliocentric distances $d_\odot$ (*left panel*) and proper-motion magnitudes, μ_{tot}, (*right panel*) for the 216 neutron stars with measured proper motion (*grey histograms*). For comparison, we also show the distances and proper-motion magnitudes for 216 simulated neutron stars (*blue histograms*) randomly sampled from a mock population generated with the fiducial parameters $\sigma_k = 265\,\mathrm{km\,s^{-1}}$ and $h_c = 0.18\,\mathrm{kpc}$ (*red histograms*). For the weighted sampling, we use the weight function $f(d_\odot) = d_\odot^{-1}\exp(-0.5 d_\odot)$ (see the text for more details).

could potentially influence a future comparison with observational data, we perform the following experiment. We first collect those neutron stars that have measured proper motions. As the main resource, we use the ATNF pulsar catalogue[3] (Manchester et al., 2005), but in some cases we provide proper motion results from more recent analyses (see Appendix B for details). We find a total of 417 neutron stars whose angular positions, proper motions in ICRS coordinates, spin periods, spin-period derivatives, *DM* values and distance estimates are reported in Table B.1. Out of these objects, we remove those stars that belong to the Magellanic Clouds, are associated with globular clusters or have a binary companion. We further select only those neutron stars with a spin-period derivative $\dot{P} > 10^{-17}$ to exclude those that have potentially been recycled. Finally, we consider only those for which an estimate of the heliocentric distance $d_\odot$ is available; in the case of radio pulsars we quote values that are derived from their respective *DMs* using the free-electron density model of Yao, Manchester, and Wang (2017) (YMW16 model hereafter). As some neutron stars have a *DM* that exceeds the maximum Galactic *DM* allowed by the YMW16 model, these cases are assigned a default distance of 25 kpc. We exclude those cases unless an alternative distance measurement is available. Applying these filters we obtain a sample of 216 Galactic, likely isolated neutron stars, whose positions and proper motions are illustrated in the upper panel in Figure 5.10.

In the left panel of Figure 5.11, the grey histogram shows the distribution of their heliocentric distances. Even if subject to some uncertainties due to imprecisions in the YMW16 model, the distance distribution peaks around 1 kpc, followed by a sharp exponential decrease. For a realistic pulsar distribution and in the absence of selection effects, we would expect the number of neutron stars to increase with distance due to an increase in the explored volume, until reaching a maximum at a distance of about 10 kpc, which comprises the region around the Galactic Centre (see the red histogram in the left panel of Figure 5.11). Thus, the shape of the

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

grey distribution in Figure 5.11, as expected, indicates that our observed sample of neutron stars with measured proper motions is incomplete in distance and subject to selection biases. By looking at the right panel of Figure 5.11, we also note that a selection bias on distance is also reflected in the distribution of the total proper motion magnitudes (grey histogram), computed as $\mu_{\text{tot}} \equiv \sqrt{\mu_{\text{RA}}^2 + \mu_{\text{DEC}}^2}$. Indeed, since the nearest stars are also characterised by larger angular proper motions, there is a tendency to detect high proper-motion stars with higher probability.

In this first empirical approach, we follow a more agnostic approach to introduce a comparable selection effect in our simulated populations. Specifically, we use a weighted random-sampling routine to select n pulsars from our mock populations, where each simulated star is assigned a weight according to a function $f(d_{\odot})$ of its heliocentric distance. This weight function has to assign larger weights to closer neutron stars in order to ensure their higher detection probabilities and has to be chosen such that we recover the observed distance distribution with sufficient accuracy. To find $f(d_{\odot})$, we simulate a mock population with the fiducial values of the kick velocity and scale height, that is $\sigma_{\text{k}} = 265\,\text{km}\,\text{s}^{-1}$ and $h_{\text{c}} = 0.18\,\text{kpc}$, respectively (red histograms in Figure 5.11). After using a given $f(d_{\odot})$ to weight the simulated neutron stars, we sample 216 mock stars and compare their distance and proper-motion distributions (shown as blue histograms in Figure 5.11) with those of the observed sample by performing two-sample Kolmogorov-Smirnov (KS) tests. After testing various functional forms, we find that $f(d_{\odot}) = d_{\odot}^{-1} \exp(-0.5d_{\odot})$ is able to reproduce the observed distributions with a good level of accuracy. More precisely, for this choice of $f(d_{\odot})$, the KS tests performed over 1000 distinct comparisons give average p-values of ~ 0.3 and ~ 0.6 for the distance and proper-motion comparison, respectively. This means that at 95% confidence level, we cannot reject the null hypothesis that the simulated and observed samples are drawn from the same underline distribution. We have also verified that for this weight function, the KS tests always provide p-values > 0.05 when comparing the observed sample with the simulated populations for reasonable values of σ_{k} and h_{c}. Only in the cases where σ_{k} and h_{c} assume extreme values (near the edges of their respective ranges) the p-values might drop below 0.05. However, these cases are associated with simulations with extreme initial conditions that are unlikely to reproduce the observations. For our basic experiment, we further make the simplified assumptions that $f(d_{\odot})$ emulates all selection biases and that it is the same for every number n of sampled neutron stars. We stress that for the purpose of this initial analysis, we do not aim to accurately constrain the selection function that affects the observed population of neutron stars. Instead, we study how the introduction of realistic selection biases will alter the predictive power of our machine learning framework. Although one might intuitively attribute the exponential factor in $f(d_{\odot})$ to scattering in the interstellar medium at large distances, the underlying nature and the precise form of the true selection function is certainly more complicated. We expect it to encompass a series of effects due to the physics of the interstellar medium and the pulsar emission itself, as well as selection effects of pulsar surveys and pulsar searches. We reserve a more accurate study disentangling the different effects that contribute to $f(d_{\odot})$ to future work.

We then analyze how the predictive power of the CNN evolves as a function of the number n of neutron stars, sampled with the above weight function $f(d_{\odot})$. To do so, we vary n from 200 to 2000 in steps of 200, and in each case re-sample the 16384 simulated populations from run S3, where both σ_{k} and h_{c} are varied. After applying a 80/20% training-validation split, we retrain the CNN on each of the
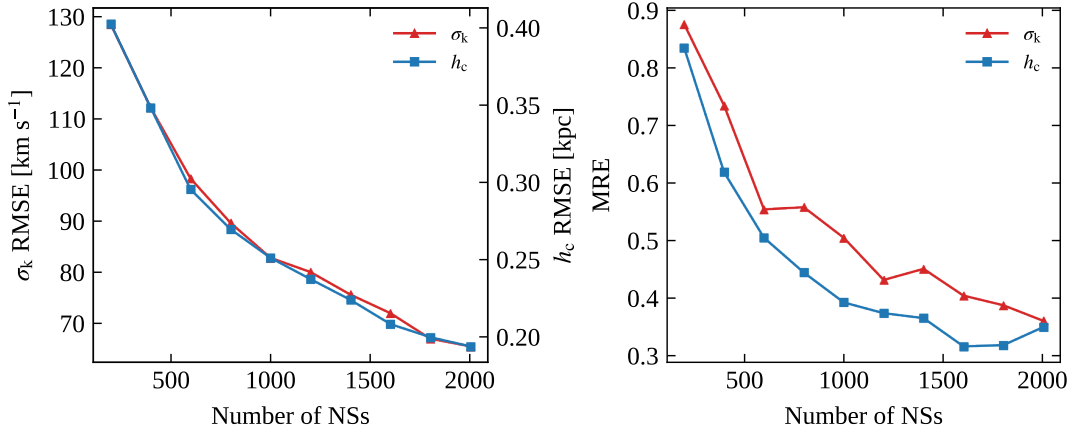
FIGURE 5.12: Trend of the RMSE (*left panel*) and MRE (*right panel*) uncertainties on the prediction of the two parameters σ_k (*blue*) and h_c (*red*) as a function of the number of neutron stars n in the resampled simulations in the validation sets. To train and validate the CNN, we use the 16384 simulated populations from simulation run S3 (with a 80/20% training/validation split), where both σ_k and h_c are varied, and which are resampled with increasing number of stars n according to the weight function $f(d_\odot) = d_\odot^{-1} \exp(-0.5 d_\odot)$.

down-sampled simulations. As before, we use the 3-channel ICRS input maps (i.e., density plus proper motion information) but instead opt for a resolution of 32×16 bins to accommodate the smaller number of stars represented in our maps. We have verified that a higher resolution of 128×64 bins does not affect the training results significantly, but slows down the training process; we therefore choose the lower resolution. We use the same training hyperparameters as in Section 5.4.1, that is an initial learning rate of 10^{-4}, a batch size of 64 and an early stop at 128 epochs. Once trained for each n value, we apply the CNN to the validation sets and compute the root-mean-square error (RMSE) and mean relative error (MRE) for the σ_k and h_c predictions as a function of n.

In the left panel of Figure 5.12, we show how the RMSE uncertainties for the predictions of the two parameters σ_k (*blue*) and h_c (*red*) diminish with increasing number of neutron stars n sampled from the simulations. We observe that both curves (with the appropriate rescaling) follow very similar trends. On the right, we show instead how the MREs evolve with n, indicating how the precision of the two-parameter prediction improves with the number of detected neutron stars. This plot shows that, under the assumptions that selection effects are unaltered and the underlying kick-velocity and birth-height distributions have a Maxwellian and exponential shape, respectively, our trained CNN is able to predict σ_k and h_c with a relative error of ~ 0.35 for a sample of 2000 stars (see also lower panel of Figure 5.10 as a reference).

These results highlight that the number of neutron stars plays a crucial role for the level of accuracy that the CNN can reach. As expected, a larger number of stars provides more information, which allows the CNN to pinpoint differences between populations evolved from different initial conditions, also when selection effects are introduced. Future observational efforts aimed at detecting and characterising new pulsars will thus play an important role in constraining the birth properties of neutron stars. Specifically, the advent of the Square Kilometer Array (SKA) will represent an important step forward into this direction. Due to its large sensitivity (a factor of 10 better than other radio telescopes) and its long baseline (up to 3000 km),

SKA has the potential to increase the number of discovered pulsars by a factor 10 (Smits et al., 2009; Smits et al., 2011). This will allow more precise timing and astrometric measurements of pulsar positions as well as distances and proper motions. A larger and more precise dataset could also help to better constrain the shape of the kick-velocity distribution and differentiate between models that try to explain the origin of the natal kicks (Tauris et al., 2015).

5.6 Summary

In this work, we have analyzed the possibility of using machine-learning (ML) techniques to reconstruct the dynamical birth properties of an evolved population of isolated neutron stars. For this purpose, we developed a simplistic population-synthesis pipeline to simulate the dynamical evolution of Galactic neutron stars and subsequently used these simulations to train and validate two different neural networks. We specifically focused on their ability to predict two parameters that strongly impact on the phenomenology of the evolved population: the dispersion, σ_k, of a Maxwell kick-velocity distribution and the scale height, h_c, of an exponential distribution for the Galactic birth heights. This was achieved by providing the networks with two-dimensional stellar density and velocity maps in galactocentric and equatorial (ICRS) coordinate frames. We found that a convolutional neural network (CNN) is able to estimate the physical parameters with high accuracy when multiple input channels, i.e., position and velocity information, are provided. In particular, when simultaneously predicting σ_k and h_c from ICRS maps, the network is able to reach absolute uncertainties lower than $10\,\mathrm{km\,s^{-1}}$ and $0.05\,\mathrm{kpc}$, respectively, corresponding to a relative error of around 10^{-2} for both parameters. Albeit obtained under simplified assumptions, our feasibility study has thus demonstrated that ML techniques are indeed suitable to infer information about the pulsar population. Our phenomenological analysis incorporating proper-motion measurements (an attempt at including observational biases in an agnostic way) has further highlighted that increasing the sample of known pulsars and accurately measuring their current characteristics with future telescopes is crucial to tightly constrain the birth properties of the neutron stars in the Milky Way. In particular, our trained CNN is able to predict σ_k and h_c with a relative error of ~ 0.35 for a sample of 2000 pulsars with measured proper motions.

We also demonstrated that one of the main factors in limiting the accuracy of the CNN's predictions in our set-up is the degeneracy between σ_k and h_c; as they both affect the evolved populations in a similar way, the network struggles to disentangle their effects. This limitation is a direct consequence of our simplified population-synthesis framework. In future works, we will go beyond modelling the dynamical evolution and focus on incorporating additional physics such as magneto-thermal and spin-period evolution. We will further model their emission in different electromagnetic bands and study corresponding detectability limits by addressing selection effects as well as observational survey biases. Such additional input information could potentially break the degeneracy between the kick-velocity and the Galactic height distributions and provide more accurate model constraints on σ_k and h_c as well as other input parameters.

The ultimate goal will be to use multi-wavelength observations of the Galactic neutron star population and take advantage of ML, combined with population synthesis, to recover their birth properties, such as the natal kick-velocity, spin-period or magnetic-field distribution. In the next chapter we will focus on inferring the

magneto-rotational properties of the pulsar population by adopting a more sophisticated simulation based inference approach.

Chapter 6

Simulation-based inference for pulsar population synthesis

6.1 Introduction

As one of the end points of stellar evolution of massive stars (see Section 1.2), neutron stars are influenced by many extremes of physics. As a result, these compact objects are connected with several of the most energetic transient phenomena in our Universe. In particular, neutron stars have been associated with fast radio bursts (e.g., Platts et al., 2019; Petroff, Hessels, and Lorimer, 2022; Zhang, 2023), superluminous supernovae (e.g., Kasen and Bildsten, 2010; Woosley, 2010; Metzger et al., 2014; Prasanna et al., 2023), ultra-luminous X-ray sources (e.g., Bachetti et al., 2014; Mushtukov et al., 2015; Fürst et al., 2016; Israel et al., 2017), long-duration (e.g., Wheeler et al., 2000; Thompson, Chang, and Quataert, 2004; Bucciantini et al., 2008; Metzger, Berger, and Margalit, 2017; Margalit et al., 2018) and short-duration (Hurley et al., 2005; Metzger, Quataert, and Thompson, 2008; Metzger and Bower, 2014; Berger, 2014; Goldstein et al., 2017) gamma-ray bursts, and detectable gravitational-wave emission (Abbott et al., 2017b; Abbott et al., 2017c; Abbott et al., 2017a; Abbott et al., 2020). Accurately modelling these processes requires a detailed understanding of neutron-star properties, which also set constraints on massive stellar evolution. Inferring the birth properties of neutron stars and the physics that govern their subsequent evolution is, thus, crucial for other fields of astrophysics.

Detecting and accurately characterizing individual objects within the entire neutron-star population is, hence, critical. However, although the number of known pulsars (those neutron stars that are observed to emit regular electromagnetic pulses) has steadily increased since the first detection in 1967 (Hewish et al., 1968), we currently only know around 3,000 of these objects (Manchester et al., 2005, see Section 1.13).[1] With around 1,100 members, isolated radio pulsars constitute the largest subclass of neutron stars (Posselt et al., 2023). However, this is only a tiny fraction of the approximately one billion neutron stars expected in our Milky Way alone.

To bridge the gap between expected and observed neutron stars, we take advantage of population synthesis (see also Section 1.15 and Chapter 5). This approach relies on producing a large catalogue of synthetic pulsar populations which are passed through a set of filters to mimic observational constraints. The resulting populations are then contrasted with the true observed sample to find those parameter regions that best explain the data. Although different versions of this methodology have been applied to pulsar data for several decades (e.g., Narayan and Ostriker, 1990; Lorimer, 2004; Faucher-Giguère and Kaspi, 2006; Gonthier et al., 2007; Bates et al., 2014; Gullón et al., 2014; Gullón et al., 2015; Cieślar, Bulik, and Osłowski, 2020),

---

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

the complexity of models that capture the properties of observed Galactic neutron stars significantly complicates the comparison between the simulated populations and the observed one. This is especially true if we are interested in quantifying uncertainties for our neutron-star parameters, because Bayesian MCMC methods (the standard tool for this kind of question, see, e.g., Feroz, Hobson, and Bridges (2009), Foreman-Mackey et al. (2013), Sharma (2017), Ashton et al. (2019), and Speagle (2020)) become infeasible for pulsar population synthesis unless significant simplifications for simulation models and the likelihood function are made (Cieślar, Bulik, and Osłowski, 2020). The main reason for this is that we can no longer write down an explicit likelihood for realistic neutron-star simulation frameworks. In this paper, we, thus, focus on simulation-based inference (SBI) (also known as likelihood-free inference; see Chapter 2 and Cranmer et al., 2020 for a recent review) in the context of pulsar population synthesis for the first time.

In the past few years, SBI has successfully challenged traditional approaches such as approximate Bayesian computation (e.g., Rubin, 1984; Beaumont, Zhang, and Balding, 2002; Dean et al., 2011; Frazier et al., 2017) in those areas of science that rely on complex simulators which lead to intractable likelihoods. The existence of such a simulator, essentially acting as a forward model, is the only requirement for SBI. As such, the approach is ideal for astrophysics and has been recently applied to parameter estimation in, e.g., cosmology (Alsing et al., 2019; Lemos et al., 2023; Lin et al., 2023; Hahn et al., 2023), high-energy astrophysics (Mishra-Sharma and Cranmer, 2022; Huppenkothen and Bachetti, 2022), gravitational-wave astronomy (Dax et al., 2021; Cheung et al., 2022; Bhardwaj et al., 2023) and exoplanet research (Vasist et al., 2023). SBI is particularly powerful in combination with neural networks, whose benefits for pulsar population synthesis studies we outlined in Ronchi et al. (2021) (see Chapter 5) by inferring point estimates for the dynamical properties of radio pulsars in the Milky Way.

In the following, we will improve on our earlier analysis and take a Bayesian perspective to infer posteriors of neutron-star parameters using SBI. For this purpose, we expand our simulator to not only model the Galactic neutron-star dynamics but also capture the magneto-rotational evolution and relevant detection limits and biases of existing radio surveys. We then construct an SBI pipeline, which we train, validate and test on a large database of synthetic P-$\dot{P}$ diagrams to infer posterior distributions of our input parameters. We specifically focus on five parameters related to the initial period distribution of pulsars and their magnetic-field properties that crucially affect the positions of stars in the P-$\dot{P}$ plane. We then apply our optimised deep-learning framework, for the first time, to the radio pulsars detected in 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) (all recorded with Murriyang, the Parkes radio telescope).

This chapter, based on the work Graber, Ronchi, Pardo-Araujo and Rea (2023) (submitted to ApJ), is structured as follows: In Section 6.2, we summarise our population synthesis framework, specifically highlighting advancements to those discussed in Chapter 5 (Ronchi et al., 2021). We then provide a general overview of SBI as well as our choice of set-up in Sections 6.3.1 and 6.3.2, respectively, while Section 6.3.3 summarises the machine-learning experiments conducted for this study. We next provide network training and inference results plus corresponding validation approaches in Section 6.4. Finally, we discuss our results and provide conclusions in Section 6.5.

6.2 Pulsar population synthesis

Our population synthesis framework broadly follows earlier works (see, e.g., Faucher-Giguère and Kaspi (2006), Gullón et al. (2014), Cieślar, Bulik, and Osłowski (2020), and Ronchi et al. (2021)). In particular, our simulator employs a Monte-Carlo approach to sample the relevant neutron-star parameters at birth from corresponding probability density functions and subsequently evolves these parameters forward in time. However, to save computation time, we do not follow the procedure in Chapter 5 (see also Ronchi et al., 2021) and evolve the dynamical properties for each simulation run. As the dynamical and magneto-rotational properties are independent, we instead simulate a single dynamical database for a large number of current pulsar positions and velocities, and subsequently sample from these distributions before determining the magneto-rotational evolution. Prescriptions for both are summarised below in Sections 6.2.1 and 6.2.2, respectively. We subsequently discuss the radio emission properties in Section 6.2.3 and then run snapshots of the total pulsar population at the current time through a set of filters to mimic observational limitations (see Section 6.2.4). The resulting simulator output are synthetic P-$\dot{P}$ diagrams of the observed pulsar population (see Section 6.2.5) that will serve as the starting point for our SBI pipeline discussed in Section 6.3.

6.2.1 Dynamical evolution

To create our dynamical database, we simulate 10^7 neutron stars from birth to today. For each object, we randomly assign an age sampled from a uniform distribution up to a maximum age of 10^8 yr, which ensures a constant birth rate of 10 stars per century. Note that this value, although somewhat higher than typically assumed (e.g., Keane and Kramer, 2008; Rozwadowska, Vissani, and Cappellaro, 2021), was chosen to populate our synthetic Milky Way with a sufficient number of neutron stars within reasonable computation time. As objects older than 10^8 yr are no longer detectable as radio pulsars (see below), the following approach provides a realistic description of the current positions of these sources.

We then define a cylindrical reference frame, (r, ϕ, z), whose origin is located at the Galactic centre. Here, r, ϕ and z denote the distance from the origin in kpc, the azimuthal angle in radians and the distance from the Galactic plane in kpc, respectively. In particular, we position our Sun at $r = 8.3$ kpc, $\phi = \pi/2$, and $z = 0.02$ kpc (see Pichardo et al., 2012, and references therein).

To determine the birth locations of individual neutron stars, we address the distributions of their massive progenitors in the (r, ϕ)-plane and along z separately. Considering the distribution of free electrons as a tracer of star formation in the Milky Way which correlates with the massive OB stars that evolve into neutron stars, we sample the initial positions in r, ϕ according to the Galactic electron density distribution of Yao, Manchester, and Wang, 2017 (see Figure 1.10). This will also allow consistency when relating pulsar distances with their dispersion measures *DM*s in Section 6.2.4. In addition, as the Galactic matter distribution is not static, we assume that the Milky Way rotates rigidly in clockwise direction with an angular velocity $\Omega = 2\pi/T$, where $T \approx 250$ Myr (Vallée, 2017; Skowron et al., 2019). For a given stellar age, we can thus retrace the angular coordinate, ϕ, at birth.

Moreover, we assume that pulsar birth positions along the z-direction follow an exponential disk model (Wainscoat et al., 1992) and sample from a probability

density function of the form

$$\mathcal{P}(z) = \frac{1}{h_c} \exp\left(-\frac{|z|}{h_c}\right). \tag{6.1}$$

We follow the pulsar population studies of Chapter 5 (see also Gullón et al., 2014; Ronchi et al., 2021) and set the characteristic scale height, h_c, to a fiducial value of 0.18 kpc. Note that this is consistent with the distribution of young, massive stars in our Galaxy (Li et al., 2019). We then randomly assign each star's z-coordinate a positive or negative sign to distribute our population above and below the Galactic plane.

Next, we focus on the pulsars' birth velocities, which are a combination of the kick velocity, v_k, imparted during the supernova due to explosion asymmetries (see Janka, Wongwathanarat, and Kramer, 2022; Coleman and Burrows, 2022, and references therein), and the velocity, v_{pr}, inherited from the progenitors' orbital Galactic motion. Specifically, we sample the magnitude of the kick velocities, $v_k \equiv |v_k|$, from a Maxwell distribution,

$$\mathcal{P}(v_k) = \sqrt{\frac{2}{\pi}} \frac{v_k^2}{\sigma_k^3} \exp\left(-\frac{v_k^2}{\sigma_k^2}\right), \tag{6.2}$$

and then assign a random direction to determine the kick along the r-, ϕ- and z-directions. For the dispersion parameter, σ_k, we take a fiducial value of $\sigma_k \approx 260\,\mathrm{km\,s^{-1}}$ (Hobbs et al., 2005), which is broadly consistent with observed proper motions of radio pulsars (Hobbs et al., 2005; Faucher-Giguère and Kaspi, 2006; Verbunt, Igoshev, and Cator, 2017; Igoshev, 2020).

The second velocity component due to the progenitors' motion depends on the Galactic gravitational potential, Φ_{MW}, and points along the azimuthal direction:

$$v_{pr} = \sqrt{r \frac{\partial \Phi_{MW}(r,z)}{\partial r}} \, \hat{\boldsymbol{\phi}}, \tag{6.3}$$

where $\hat{\boldsymbol{\phi}}$ is a unit vector in ϕ-direction. For this study, we consider a Galactic potential that is given as the sum of four components, i.e., the nucleus, Φ_n, the bulge, Φ_b, the disk, Φ_d, and the halo, Φ_h, (Marchetti, Rossi, and Brown, 2019). The nucleus and bulge contributions are described by a spherical Hernquist potential (Hernquist, 1990):

$$\Phi_{n,b} = -\frac{GM_{n,b}}{R_{n,b} + R}, \tag{6.4}$$

where $R = \sqrt{r^2 + z^2}$ is the spherical radial coordinate and G the gravitational constant. The disk has a cylindrical Miyamoto-Nagai potential of the form (Miyamoto and Nagai, 1975)

$$\Phi_d = -\frac{GM_d}{\sqrt{\left(a_d + \sqrt{z^2 + b_d^2}\right)^2 + r^2}}, \tag{6.5}$$

where a_d and b_d represent the scale length and scale height of the disk, respectively. Finally, the halo is characterised by a spherical Navarro-Frenk-White potential (Navarro, Frenk, and White, 1996):

$$\Phi_h = -\frac{GM_h}{R} \ln\left(1 + \frac{R}{R_h}\right). \tag{6.6}$$

The free parameters, $M_{\text{n,b,d,h}}$, $R_{\text{n,b,h}}$, a_{d} and b_{d}, can be obtained through fits of the Milky Way's mass profile and are given in Table 5.2 (see also Tab. 1 of Marchetti, Rossi, and Brown (2019) and Bovy (2015)).

After determining the initial positions and velocities for each of our 10^7 neutron stars, we perform the dynamical evolution by solving the Newtonian equation of motion in cylindrical coordinates, $\ddot{\boldsymbol{r}} = -\boldsymbol{\nabla}\Phi_{\text{MW}}$, according to the stars' respective ages. In this way, we obtain a database of current pulsar positions and velocities in the Milky Way.

6.2.2 Magneto-rotational evolution

The primary diagnostic for the pulsar population is the P-$\dot{P}$ diagram. For our population synthesis study, we focus on rotation-powered radio pulsars, the largest class of neutron stars (see Section 1.13). Corresponding period and period-derivative measurements for this population are enabled via radio timing. To first order, radio pulsars can be approximated as rotating magnetic dipoles, implying that their spin-down is driven by electromagnetic dipole radiation (see Section 1.8). The locations of individual neutron stars, and the shape of the population's distribution, in the P-$\dot{P}$ plane are, hence, determined by their dipolar magnetic fields and rotation periods at birth, and the subsequent magneto-rotational evolution. The latter couples the evolution of the pulsar period, P, the dipolar magnetic-field strength, B, at the pole and the inclination angle, χ, between the magnetic and the rotation axis.

To capture these physics, we first sample the misalignment angle at birth, χ_0, randomly in the range $[0, \pi/2]$ according to the probability density (Gullón et al., 2014)

$$\mathcal{P}(\chi_0) = \sin \chi_0. \tag{6.7}$$

We then sample the logarithm of the initial period, P_0, (measured in s) and the initial magnetic field, B_0, (measured in G) for each pulsar from normal distributions of the form (Popov et al., 2010; Gullón et al., 2014; Igoshev, 2020; Igoshev et al., 2022; Xu et al., 2023)

$$\mathcal{P}(\log P_0) = \frac{1}{\sqrt{2\pi}\sigma_{\log P}} \exp\left(-\frac{\log P_0 - \mu_{\log P}}{2\sigma_{\log P}^2}\right), \tag{6.8}$$

$$\mathcal{P}(\log B_0) = \frac{1}{\sqrt{2\pi}\sigma_{\log B}} \exp\left(-\frac{\log B_0 - \mu_{\log B}}{2\sigma_{\log B}^2}\right). \tag{6.9}$$

The means, $\mu_{\log P}$, $\mu_{\log B}$, and the standard deviations, $\sigma_{\log P}$, $\sigma_{\log B}$, are free parameters of our model and four of those parameters, whose posteriors we set out to infer with our SBI approach in Section 6.3. We will specifically explore the ranges $\mu_{\log P} \in [-1.5, -0.3]$, $\mu_{\log B} \in [12, 14]$, $\sigma_{\log P} \in [0.1, 1.0]$ and $\sigma_{\log B} \in [0.1, 1.0]$ to encompass results of earlier analyses (e.g., Gullón et al., 2014).

Assuming that pulsars spin down due to dipolar emission, we follow Philippov, Tchekhovskoy, and Li (2014) and Spitkovsky (2006) and solve the following coupled
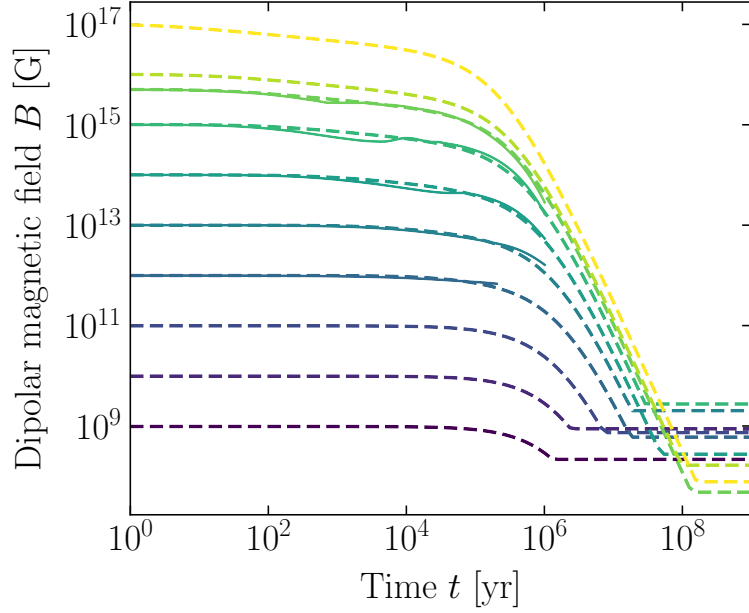
FIGURE 6.1: Illustration of the *B*-field parametrisation used for this study. The five solid curves represent realistic two-dimensional simulations of magneto-thermal evolution in the neutron-star crust (Viganò et al., 2021). We fit these together with the late-time power-law evolution of the magnetic field with several broken power laws. The dashed curves shown here are determined for $a_\text{late} = -3.0$. The colours represent the initial magnetic-field strength, B_0. To avoid the field decaying to unrealistically small numbers at very late times, we sample the final fields from a Gaussian distribution. The procedure, which allows us to easily extract the dipolar field strength, B, at different times, t, to study the magneto-rotational evolution of our synthetic pulsars, is described in details in Appendix C.

differential equations (see Equations (1.42)):

$$\dot{P} = \frac{\pi^2}{c^3} \frac{B^2 R_\text{NS}^6}{I_\text{NS} P} \left(\kappa_0 + \kappa_1 \sin^2 \chi \right), \tag{6.10}$$

$$\dot{\chi} = -\frac{\pi^2}{c^3} \frac{B^2 R_\text{NS}^6}{I_\text{NS} P^2} \kappa_2 \sin \chi \cos \chi, \tag{6.11}$$

where c is the speed of light, $R_\text{NS} \approx 11\,\text{km}$ the neutron-star radius and $I_\text{NS} \simeq 2M_\text{NS}R_\text{NS}^2/5 \approx 1.36 \times 10^{45}\,\text{g}\,\text{cm}^2$ the stellar moment of inertia (for a fiducial mass $M_\text{NS} \approx 1.4\,\text{M}_\odot$). For realistic pulsars surrounded by plasma-filled magnetospheres, we choose $\kappa_0 \simeq \kappa_1 \simeq \kappa_2 \simeq 1$, and note that Eqn. (6.11) implies that χ decreases with time, i.e., our pulsars move towards alignment.

The final ingredient is a suitable prescription for the evolution of the dipolar magnetic-field strength. While the *B*-field decay in the neutron-star crust is typically assumed to be driven by the combined action of the Hall effect and Ohmic dissipation (e.g., Aguilera, Pons, and Miralles, 2008a), changes in the magnetic field are strongly coupled to the thermal properties of the neutron-star interior (e.g., Pons and Viganò, 2019). This is particularly important for strongly magnetised neutron stars with fields above $\sim 10^{13}\,\text{G}$ and, hence, relevant for a significant fraction of our simulated pulsar population. In the past decade, several theoretical and numerical

efforts have begun to unveil the complex processes of magneto-thermal evolution (e.g., Viganò et al., 2013; Viganò et al., 2021; De Grandis et al., 2021; Igoshev et al., 2021; Dehman et al., 2023, see also Section 1.6). As corresponding simulations are highly time-consuming, we instead develop a new approach, outlined in detail in Appendix C and summarised in Figure 6.1, that parametrises a range of magneto-thermal simulations for different magnetic-field strengths (Viganò et al., 2021). This prescription allows us to extract magnetic fields up to pulsar ages of around 10^6 yr. Above this value, current numerical simulations become unreliable because they rely on implementations of complex microphysics that are unsuitable for cold, old stars. Moreover, they do not capture the highly uncertain physics of neutron-star cores, which become relevant at large ages. We instead incorporate the field evolution at late times by means of a power law of the form

$$B(t) \propto \left(1 + \frac{t}{\tau_{\text{late}}}\right)^{a_{\text{late}}},\tag{6.12}$$

where $\tau_{\text{late}} \approx 2 \times 10^6$ yr, t is the time, and the power-law index, a_{late}, is the fifth free parameter of our model. We note that although the details of core field evolution are not known, Equation (6.12) is physically motivated because several known mechanisms exhibit similar power-law behaviour (see Appendix C). We will, hence, explore the parameter range $a_{\text{late}} \in [-3.0, -0.5]$. Finally, to prevent the dipolar magnetic field from decaying to arbitrarily small values (in disagreement with observations of old, recycled millisecond pulsars; see, e.g., Lorimer (2008)), we assume that the field eventually reaches a constant value. Therefore, we sample the logarithm of the field, B_{final}, from a normal distribution with mean $\mu_{\log B,\text{final}} = 8.5$ and a standard deviation $\sigma_{\log B,\text{final}} = 0.5$ in line with observations of old pulsars.

Following this prescription allows us to determine the spin periods, dipolar field strengths and misalignment angles for our simulated pulsars at the current time.

6.2.3 Emission characteristics

We next implement a prescription for the radio emission geometry to determine those pulsars whose beams sweep over the Earth and are, in principle, detectable. In the *canonical* model of radio pulsars, their emission is produced close to the stellar surface in the cone-shaped, open field-line region (Lorimer and Kramer, 2012; Johnston et al., 2020, see also Section 1.9). Assuming that this entire region is involved in the emission, geometric considerations allow us to estimate the half opening angle of the emission beam, ρ_{b}, (in rad) via (see Equation (1.57) and Gangadhara and Gupta, 2001)

$$\rho_{\text{b}} \simeq \sqrt{\frac{9\pi r_{\text{em}}}{2cP}},\tag{6.13}$$

where r_{em} is the emission height. The latter is thought to be period independent and we set it to 300 km following Johnston et al. (2020) (see also references therein). Note that several studies of pulsars with stable emission profiles have recovered this $\rho_{\text{b}} \propto P^{-1/2}$ behaviour (e.g., Kramer et al., 1994; Maciesiak and Gil, 2011; Skrzypczak et al., 2018, and see also Figure 1.7). Knowledge of ρ_{b}, then, allows us to obtain the solid angle, Ω_{b}, covered by a pulsar's two radio beams (see Equation (1.78)). More specifically,

$$\Omega_{\text{b}} = 4\pi(1 - \cos\rho_{\text{b}}).\tag{6.14}$$

As we do not expect biases in how we observe this conal emission for any given pulsar, we draw a random line-of-sight angle, α, with respect to the rotation axis in the range $[0, \pi/2]$ using the probability density $\sin \alpha$. Combined with the half opening angle, ρ_b, and the evolved inclination angle, χ, we can then determine those pulsars whose radio beams are visible from Earth. We note that as a result of this purely geometric argument, between $\sim 60 - 95\%$ of our generated pulsars (depending on the specific choice of magneto-rotational parameters) are typically not detectable.

We proceed with determining the emission characteristics of those neutron stars that point towards the Earth. In particular, we follow Maciesiak, Gil, and Ribeiro (2011) and express the intrinsic pulse width (measured in s) of our simulated pulsars as follows (see Equation (1.61)):

$$w_{\text{int}} = \frac{2}{\pi} \arcsin \sqrt{\frac{\sin^2\left(\frac{\rho_b}{2}\right) - \sin^2\left(\frac{\alpha-\chi}{2}\right)}{\sin(\alpha)\sin(\chi)}} \, P, \tag{6.15}$$

where we replace $\beta = \alpha - \chi$. Finally, as the radio emission is ultimately powered by the stars' rotational energy reservoir, we follow the common procedure (e.g., Faucher-Giguère and Kaspi, 2006) and assume that the intrinsic radio luminosity, L_{int}, (in erg s^{-1}) for each pulsar depends on the spin-down power, $\dot{E}_{\text{rot}} \propto \dot{P}/P^3$ (see Section 1.9.4). In particular, we assume

$$L_{\text{int}} = L_0 \sqrt{\frac{\dot{P}}{P^3}}. \tag{6.16}$$

where L_0 is a normalisation factor whose logarithm we sample from a normal distribution with mean $\mu_{\log L} = 35.5$ and standard deviation $\sigma_{\log L} = 0.8$ to eventually recover observed luminosities (see also Faucher-Giguère and Kaspi, 2006; Gullón et al., 2014).

6.2.4   Simulating detections

Armed with the knowledge of intrinsic pulsar properties, we now turn to the possibility of detecting those objects whose emission beams cross our line of sight. Following Section 1.10, first, the bolometric radio flux, S, that reaches us from any given simulated pulsar is equal to

$$S = \frac{L_{\text{int}}}{\Omega_b d^2}, \tag{6.17}$$

where d is the distance known from the dynamical evolution outlined in Section 6.2.1. To determine the corresponding radio flux density, S_f, (measured in Jy) at a specific observing frequency, f, we follow Lorimer and Kramer (2012) and assume that the radio emission spectrum follows a power law in f (see Equation (1.79)). In particular, we set the spectral index to -1.6 (Jankowski et al., 2018). We can, hence, approximate the total fluence of a pulse with width, w_{int}, as $S_f w_{\text{int}}$. Assuming that this fluence stays constant as the radio signal propagates from the pulsar to us, we estimate the flux density, $S_{f,\text{obs}}$, that reaches Earth as

$$S_{f,\text{obs}} \simeq S_f \frac{w_{\text{int}}}{w_{\text{obs}}} \tag{6.18}$$

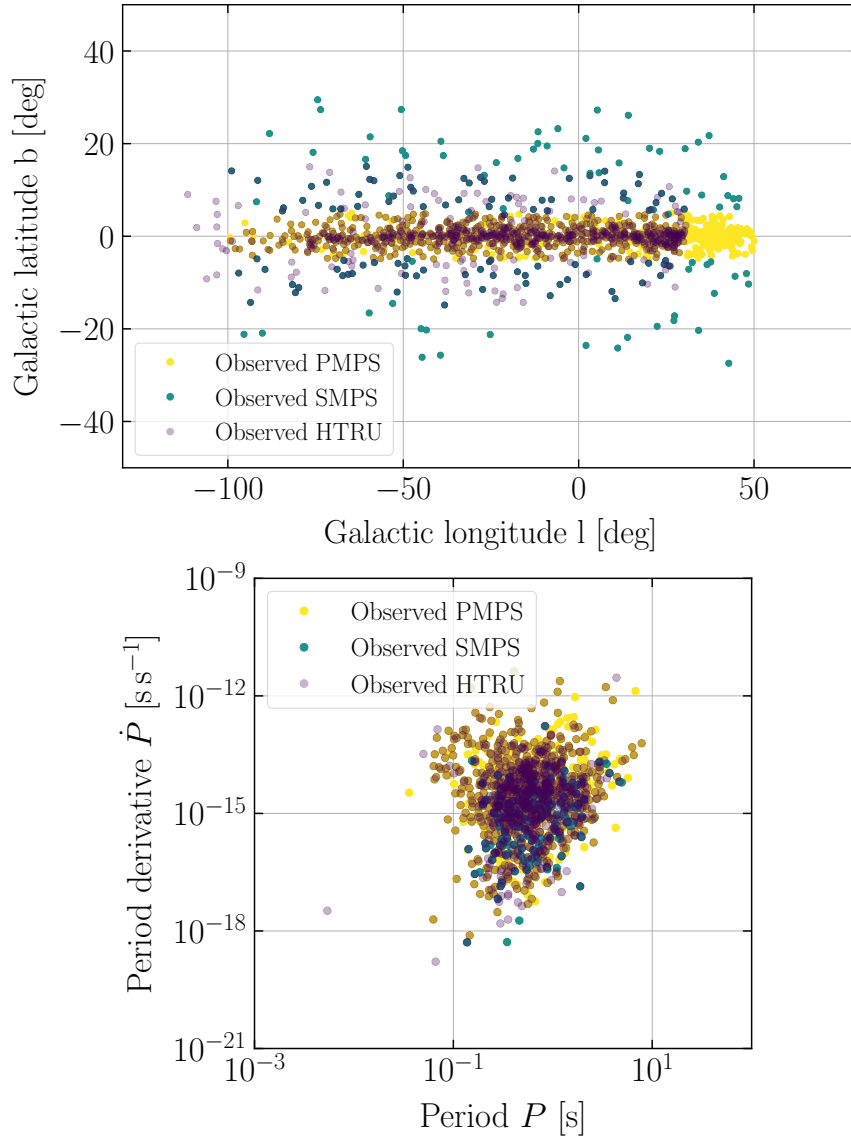where w_{obs} is the observed pulse width.

FIGURE 6.2: Observed populations of isolated Galactic radio pulsars detected with the Parkes Multibeam Pulsar Survey (PMPS), the Swinburne Intermediate-latitude Pulsar Survey (SMPS) and the low- and mid-latitude High Time Resolution Universe (HTRU) surveys (highlighted in yellow, light blue and purple, respectively). The top panel shows the distribution of these populations in Galactic latitude, b, and longitude, l, while the bottom panel depicts the detected pulsars in the period, P, and period derivative, $\dot{P}$, plane. Data taken from the ATNF Pulsar Catalogue (Manchester et al., 2005, https://www.atnf.csiro.au/research/pulsar/psrcat/). Observational filters are described in detail in the text.

TABLE 6.1: Survey parameters for the Parkes Multibeam Pulsar Survey (PMPS), the Swinburne Intermediate-latitude Pulsar Survey (SMPS), the low- and mid-latitude High Time Resolution Universe (HTRU) surveys taken from Manchester et al. (2001) and Lorimer et al. (2006), Edwards et al. (2001) and Jacoby et al. (2009) and Keith et al. (2010), respectively. We provide the survey region in Galactic longitude, l, and latitude, b, the central observing frequency f, the channel width Δf, the sampling time τ_{samp}, the telescope gain, G, the number of observed polarisations, n_{pol}, the observing bandwidth Δf_{bw}, the integration time t_{obs}, the degradation factor β, the system temperature, T_{sys}, and the S/N threshold for each of the surveys. Corresponding units are given in brackets in the first column.

Survey	PMPS	SMPS	HTRU mid	HTRU low
sky region	$-100° < l < 50°$	$-100° < l < 50°$	$-120° < l < 30°$	$-80° < l < 30°$
	$\|b\| < 5°$	$5° < \|b\| < 30°$	$\|b\| < 15°$	$\|b\| < 3.5°$
f (GHz)	1.374	1.374	1.352	1.352
Δf (kHz)	3000	3000	390.625	390.625
τ_{samp} (μs)	250	125	64	64
G (K Jy^{-1})	0.735	0.735	0.735	0.735
n_{pol}	2	2	2	2
Δf_{bw} (MHz)	288	288	340	340
t_{obs} (s)	2100	265	540	4300
β	1.5	1.5	1.5	1.5
T_{sys} (K)	21	21	23	23
S/N threshold	9	9	9	9

Specifically, as a radio pulse propagates, it experiences dispersion and scattering caused by interactions with the free electrons and density fluctuations in the interstellar medium (ISM), respectively. Both mechanisms result in a broader pulse when compared with the intrinsic width, w_{int}. Further broadening is caused by instrumental effects, which are dominated by the sampling time, τ_{samp}, of the hardware used to record radio observations (see Section 1.10). Accounting for these processes, we can write the observed pulse width as (Cordes and McLaughlin, 2003):

$$w_{\text{obs}} \simeq \sqrt{w_{\text{int}}^2 + \tau_{\text{samp}}^2 + \tau_{DM}^2 + \tau_{\text{scat}}^2}. \tag{6.19}$$

We follow Bates et al. (2014) to determine τ_{DM}, encoding the pulse smearing due to dispersion for a single frequency channel of the telescope's receiver (see Equation 1.86). Specifically,

$$\tau_{DM} = \frac{e^2}{\pi m_e c} \frac{\Delta f_{\text{ch}}}{f^3} DM, \tag{6.20}$$

where e is the electronic charge, m_e the corresponding mass, Δf the hardware-specific width of a frequency channel at observing frequency, f, and DM is the dispersion measure. We further use the empirical fit relationship from Krishnakumar et al. (2015) for τ_{scat}, the pulse smearing due to scattering of radio waves by an inhomogeneous and turbulent ISM:

$$\tau_{\text{sc}} = 3.6 \times 10^{-9}\,\text{s} \left(\frac{DM}{1\,\text{pc cm}^{-3}} \right)^{2.2} \left[1.0 + 1.94 \times 10^{-3} \left(\frac{DM}{1\,\text{pc cm}^{-3}} \right)^{2.0} \right]. \tag{6.21}$$

where τ_{scat} is measured in s. We moreover account for a significant scatter in the

underlying data (see Fig. 3 in Krishnakumar et al., 2015) by drawing $\log \tau_{\text{scat}}$ values from a Gaussian distribution around the fit (6.21) with a standard deviation of 0.5. We also incorporate the fact that Krishnakumar et al. (2015) analysed observations at 327 MHz by rescaling to a given observing frequency, f, assuming a Kolmogorov spectrum, i.e., $\tau_{\text{scat}} \propto f^{-4.4}$ (see Lorimer and Kramer, 2012, for details). As both τ_{DM} and τ_{scat} depend on the pulsars' respective dispersion measure, we again employ the Galactic electron density distribution of Yao, Manchester, and Wang, 2017 to convert our simulated neutron-star positions from Sec. 6.2.1 into DM values.

At this stage, we require information for the radio surveys we want to emulate. We specifically focus on three surveys recorded 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). All relevant survey parameters (including the sampling time, τ_{samp}, the observing frequency, f, and the channel width, Δf, needed to calculate w_{obs}) are summarised in Table 6.1.

To assess if those simulated stars that cross our line of sight are detectable with a given survey, we first determine if they are located in the surveys' fields of view. While PMPS and HTRU have a similar sky coverage, we highlight that SMPS detected pulsars at higher Galactic latitude (see left panel of Figure 6.2). This survey is, thus, sensitive to older neutron stars which have had sufficient time to move away from their birth positions closer to the Galactic plane, providing complementary information on the pulsar population. For those objects that fall within our survey coverage, we subsequently establish if they are sufficiently bright to be detected. To do so, we calculate the pulsars' signal to noise ratio, S/N, using the radiometer equation (see Section 1.12 and Lorimer and Kramer, 2012):

$$S/N = \frac{S_{f,\text{mean}} G \sqrt{n_{\text{pol}} \Delta f_{\text{bw}} t_{\text{obs}}}}{\beta \left[T_{\text{sys}} + T_{\text{sky}}(l, b) \right]} \sqrt{\frac{P - w_{\text{obs}}}{w_{\text{obs}}}}. \tag{6.22}$$

Here, $S_{f,\text{mean}} \simeq S_{f,\text{obs}} w_{\text{obs}} / P$ denotes the mean flux density averaged over a single rotation period P, G is the receiver gain (see Lorimer et al., 1993; Bates et al., 2014, for details), n_{pol} is the number of detected polarisations, Δf_{bw} is the observing bandwidth, t_{obs} the integration time and $\beta > 1$ a degradation factor that accounts for imperfections during the digitisation of the signal. Moreover, T_{sys} denotes the system temperature and $T_{\text{sky}}(l, b)$ is the sky background temperature dominated by synchrotron emission of Galactic electrons which varies strongly with latitude, l, and longitude, b. To model the latter, we use results from Remazeilles et al. (2015), who provided a refined version of the temperature map of Haslam et al. (1981) and Haslam et al. (1982) (see Figure 1.11. As the underlying data were obtained at 408 MHz, we rescale to the relevant observing frequencies by assuming a power-law dependence of the form $T_{\text{sky}} \propto f^{-2.6}$ (Lawson et al., 1987; Johnston et al., 1992).

A synthetic pulsar counts as detected, if the value obtained from Equation (6.22) exceeds the surveys' sensitivity thresholds. We aim to recover the numbers of detected isolated Galactic radio pulsars for each survey, i.e.

$$\begin{array}{l} \text{PMPS: 1009 observed pulsars,} \\ \text{SMPS: 218 observed pulsars,} \\ \text{HTRU: 1023 observed pulsars.} \end{array} \tag{6.23}$$

To obtain these values, we used the data from the ATNF Pulsar Catalogue (Manchester et al., 2005)[2] and removed extragalactic sources and those in globular clusters. We further applied a cut-off in period ($P > 0.01\,$s) and period derivative ($\dot{P} > 10^{-19}\,$s s^{-1}; for those pulsars with measured $\dot{P}$ values) to remove those objects that have (likely) undergone recycling from a companion star and cannot be modeled with the framework discussed so far. The locations of the observed objects in the P-$\dot{P}$ plane are shown in the right panel of Figure 6.2.

6.2.5   Simulation output

To simulate our *mock* observed pulsar populations, we do not make any assumptions on the neutron-star birth rate at this stage. Instead, we randomly sample a subset of 10^5 neutron stars from our dynamical database (see Section 6.2.1). We subsequently evolve these stars magneto-rotationally as outlined in Section 6.2.2 and assess how many of them are detected by each of the three surveys (see Sections 6.2.3 and 6.2.4), saving their respective properties. We iterate this process until the number of detected stars matches the number of observed stars in each survey. Note that we adaptively reduce the number of stars we draw from our dynamical database to 10^4 and 5×10^3, once we have recovered 90% and 95% of the target values, respectively. The output of a single simulator run, which has a typical computation time of around $1 - 2\,$hr, is a data frame containing the properties of those pulsars we can detect with PMPS, SMPS and HTRU.

The location of the resulting synthetic population and the shape of the stars' distribution in the P-$\dot{P}$ plane is directly controlled by the magneto-rotational parameters, $\mu_{\log B}, \sigma_{\log B}, \mu_{\log P}, \sigma_{\log P}$ and a_{late}, the five parameters we set out to infer in the following. Three examples of synthetic P-$\dot{P}$ diagrams are shown in the top three panels of Figure 6.3.

We note that our prescription does not require a *by-hand* implementation of a pulsar death line (e.g., Chen and Ruderman, 1993; Rudak and Ritter, 1994; Zhang, Harding, and Muslimov, 2000, see also Section 1.9.3), beyond which radio emission ceases. Instead, pulsars become naturally undetectable if they approach the bottom right of the P-$\dot{P}$ plane in our framework. This is due the evolution towards (*i*) smaller misalignment angles, χ, resulting in smaller beaming fractions, and (*ii*) smaller $\dot{P}$ and thus E_{rot}, ultimately leading to sources that are too faint to be detected.

At this point, we also highlight that our approach provides information on the number of total stars generated over a time scale of $10^8\,$yr (the oldest possible age for stars in our dynamical database), implying that we can directly determine the birth rate required to reproduce observations for a given survey. Although not the primary focus of this work, we note two things here: First, the number of detectable neutron stars per iteration step described above and, thus, the birth rate (as well as the distribution of stars in the P-$\dot{P}$ plane) depends strongly on the five magneto-rotational parameters. For some parameter combinations, the birth rate can become unrealistically large and the computation time extensive. To overcome this issue, we stop our iterative simulation approach once the birth rate exceeds a conservative limit of 5 neutron stars per century (Keane and Kramer, 2008; Rozwadowska, Vissani, and Cappellaro, 2021, see also Section 1.14). While this implies that we do not reach the numbers of observed objects in these simulations, we still use them in the following to assess if our inference approach can identify those parameter combinations as unreasonable from the distribution of stars in the P-$\dot{P}$ plane alone.

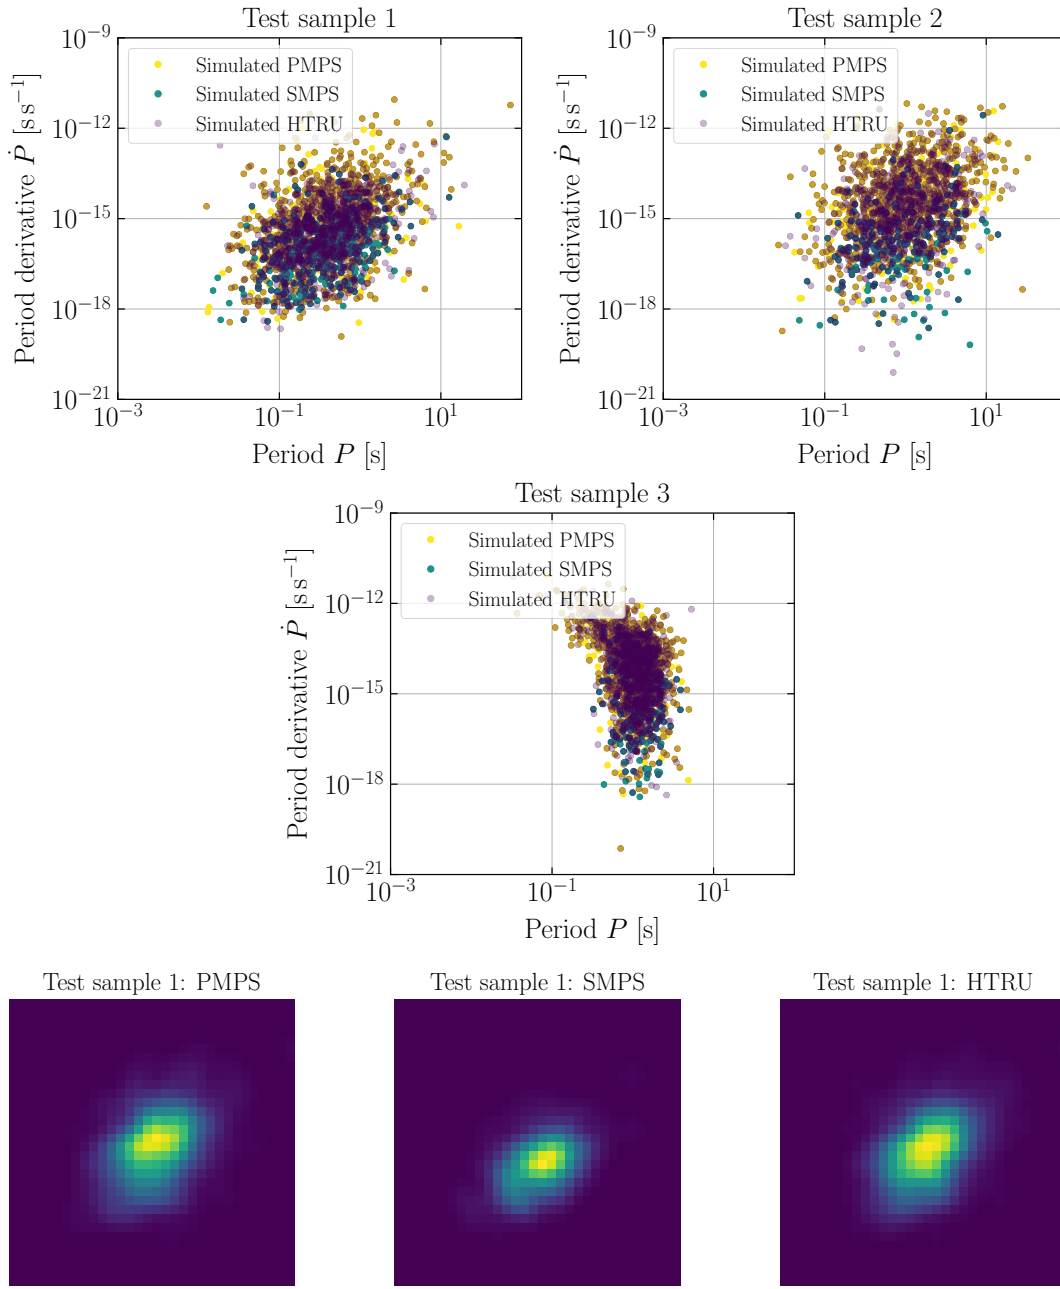[2]https://www.atnf.csiro.au/research/pulsar/psrcat/

FIGURE 6.3: Examples of simulated pulsar populations and the corresponding density maps, which are fed into the simulation-based inference pipeline. The top three panels show synthetic P-$\dot{P}$ diagrams for the three surveys considered in this study generated from three random sets of magneto-rotation parameters. In particular, test sample 1 (*top left*) is the result of a simulation with $\mu_{\log B} \approx 13.19, \sigma_{\log B} \approx 0.96, \mu_{\log P} \approx -0.85, \sigma_{\log P} \approx 0.51$ and $a_{\text{late}} \approx -0.86$, while test sample 2 (*top right*) was generated with $\mu_{\log B} \approx 13.86, \sigma_{\log B} \approx 0.88, \mu_{\log P} \approx -0.42, \sigma_{\log P} \approx 0.61$ and $a_{\text{late}} \approx -1.71$. Finally, test sample 3 (*top middle*) corresponds to $\mu_{\log B} \approx 13.35, \sigma_{\log B} \approx 0.24, \mu_{\log P} \approx -1.25, \sigma_{\log P} \approx 0.60$ and $a_{\text{late}} \approx -2.38$. The bottom row show the three density maps (one for each survey) generated with a resolution of 32 from the P-$\dot{P}$ diagram for test sample 1. Here, dark blue encodes regions where no neutron stars are present, while yellow bins represent the largest density for the binned pulsar distribution.

Second, for a single simulation run, we generally do not obtain the same birth rate for all three surveys and estimates can differ by a factor of $\sim 1 - 3$ neutron stars per century. In principle, we only expect the *correct* physical simulator to produce the observed distributions of pulsars across different surveys. The correct simulation framework is, however, not known and constraining the relevant physics is the main goal of our analysis. To explore this behaviour, we therefore produce neutron stars until the target values in all three surveys are reached (or exceeded). While this implies that the number of detected objects in some simulations can be larger than the observed number of stars for a given survey (by up to a factor of ~ 3), our focus on the location and shape of the pulsar distribution in P and $\dot{P}$ outlined below circumvents this issue. We will, however, return to the issue of the birth rate in the discussion in Section 6.5, once we have explained our inference approach and provided results for our best estimates.

To provide a broad range of synthetic P-$\dot{P}$ diagrams for our inference pipeline, we explore the ranges outlined in Section 6.2.2 and uniformly sample random combinations of the five parameters as follows:

$$
\begin{aligned}
\mu_{\log B} &\in \mathcal{U}(12, 14), \\
\sigma_{\log B} &\in \mathcal{U}(0.1, 1), \\
\mu_{\log P} &\in \mathcal{U}(-1.5, -0.3), \\
\sigma_{\log P} &\in \mathcal{U}(0.1, 1), \\
a_{\text{late}} &\in \mathcal{U}(-3, -0.5).
\end{aligned}
\tag{6.24}
$$

We generate a total of 360,000 parameter combinations (which we refer to as our input parameters or labels in the following) and simulate the corresponding synthetic populations in parallel over the course of 6 weeks.

To represent the discrete output of our simulator in a way that can be processed by a neural network, we follow our earlier study (see Chapter 5 and Ronchi et al., 2021) and convert a single P-$\dot{P}$ diagram for three surveys as seen in the top row of Figure 6.3 into three two-dimensional density maps (one for each survey) by counting the number of stars within a given bin. In particular, we set the limits $P \in [0.001, 100]\,\text{s}$ and $\dot{P} \in [10^{-21}, 10^{-9}]\,\text{s}\,\text{s}^{-1}$ and test the inference procedure for a resolution of 32 and 64 bins. To avoid sharp edges in our binned distributions, we apply a smoothing Gaussian filter (with radius 4σ and $\sigma = 1$), which will also improve the stability during the training of our machine-learning pipeline. An example of the resulting density maps is shown in the bottom row of Figure 6.3 for one of our simulations.

The final preprocessing stage for our simulated data is either a normalisation or a standardisation step (depending on the choice of set-up discussed below) to provide the neural network with signals and labels of similar magnitude (see Section 2.3). In the former case, the bins in each individual density map are rescaled such that they contain continuous values between 0 and 1. The same holds for the corresponding labels, which are normalised over the entire parameter ranges given in Equation (6.24). On the other hand, standardisation is achieved by using *z*-scores, so that the resulting information in each map has a mean of 0 and standard deviation of 1. The same method is applied to the labels across our entire set of simulations.

6.3 Simulation-based inference

6.3.1 Overview

In Section 2.4 we introduced the simulation-based inference (SBI) framework, here we will overview the main concepts. The pulsar population-synthesis pipeline summarised in Section 6.2 is a typical example of a stochastic forward model which aims to emulate real-world observations. We specifically introduced stochasticity by sampling relevant variables from underlying probability distributions using Monte-Carlo techniques. In particular, given the input parameter, $\theta = \{\theta_1, \theta_2, \dots\}$, our simulator generates a synthetic realisation of the observed data, x. The key challenge is then to constrain our model parameters in such a way that they are consistent with true observations, x_{obs}, and our prior knowledge, encoded in the prior distribution, $\pi(\theta)$. To this end, we want to compute the posterior distribution, $\mathcal{P}(\theta|x)$, using Bayes' theorem

$$\mathcal{P}(\theta|x) = \frac{\pi(\theta)\mathcal{P}(x|\theta)}{\mathcal{P}(x)}, \tag{6.25}$$

where $\mathcal{P}(x|\theta)$ is the likelihood of our data, x, given the parameter, θ, and $\mathcal{P}(x)$ denotes the evidence obtained by marginalizing the likelihood over all θ. However, for complex simulators like ours, we typically cannot write down an explicit form of the likelihood function so $\mathcal{P}(x|\theta)$ is essentially intractable. In addition, even if the likelihood were tractable, computing the evidence can become very costly as it involves complicated integrals over a multidimensional parameter space.

Simulation based inference (SBI) circumvents these issues by taking advantage of the fact that our simulator encodes the likelihood function implicitly (see Section 2.4 and Cranmer et al., 2020, for a recent review). These approaches have been particularly successful in combination with deep-learning techniques because neural networks can be used to learn a probabilistic association between a given simulation outcome, x, and the input parameters, θ. This allows an approximation of the posterior distribution, $\mathcal{P}(\theta|x)$, without the need to explicitly compute the likelihood.

For the following study, we choose the neural posterior estimation (NPE) approach (see Section 2.4.1) to directly learn the posterior conditional on our simulated data (avoiding the additional sampling step required for neural likelihood estimation (NLE) and neural ratio estimation (NRE) as seen in Section 2.4) and take advantage of the corresponding implementation in the open-source Python package `sbi` (Tejero-Cantero et al., 2020).[3]

6.3.2 Deep-learning set-up

As described in Section 2.4.1, for NPE, we approximate the posterior using a family of densities, q_ψ, characterised by the distribution parameters, ψ. For our SBI pipeline, we then use a neural network, F, to learn these ψ for our simulator output, x, by adjusting the network weights, w. In particular, we aim to optimise the neural density estimator such that $q_{F(x,w)}(\theta) \approx \mathcal{P}(\theta|x)$. As described in Section 2.4.1 this can be achieved by minimizing the expectation value of the following loss function (Papamakarios and Murray, 2016)

$$\mathcal{L}(w) = -\sum_{i=1}^{N} \log q_{F(x_i, w)}(\theta_i) \tag{6.26}$$
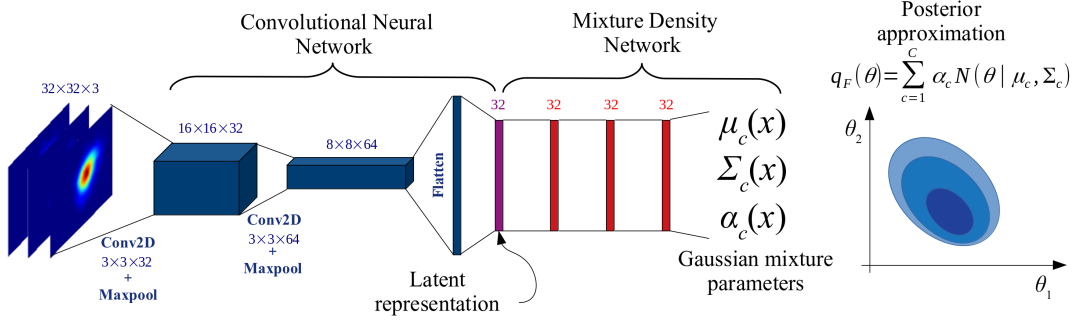
---

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

FIGURE 6.4: Schematic representation of our inference pipeline for three input $P - \dot{P}$ maps (one for each survey) with resolution 32×32. A convolutional neural network (CNN) is first used to extract features from our images and produce a compressed representation of our simulation output, x. We then train a Gaussian mixture density network (MDN), a flexible neural density estimator, on this latent representation to approximate the posterior distribution of the simulation parameters, θ.

over a training data set $\{\theta_i, x_i\}$ of size N, provided that N is large and the density estimator sufficiently flexible. In practice, we maximise the negative of $\mathcal{L}(w)$, i.e., the total log-posterior. A key advantage of the resulting posterior approximation is that the evaluation of $q_{F(x,w)}(\theta)$ corresponds to a simple forward pass through a neural network (without the need to simulate additional data), which is very fast. We will take advantage of this *amortised nature* of the posterior to assess the quality of our inferences below.

For our pulsar study, we have drawn the model parameters $\theta_i = \{\mu_{\log B}, \sigma_{\log B}, \mu_{\log P}, \sigma_{\log P}, a_{\text{late}}\}$ from uniform priors as defined previously in Equation (6.24). The corresponding output, x_i, of a single run through the simulator are the three P-$\dot{P}$ density maps (one for each survey) illustrated in the bottom row of Fig. 6.3. In the following, we stack these maps together to form a three-channel input for our neural network. Of the 360,000 synthetic simulations produced, we use 90% for training and validation reserving the remaining 10% for testing purposes. The former data set is further split into 90% for training (291, 600 populations) and 10% for validation (32, 400 populations). We note that as each population is represented by three density maps, we train the following inference pipeline on roughly 875, 000 images. Performance results for the unseen test samples quoted in the following are computed for 10% of the full test set (3, 600 populations) for computational reasons. The full workflow and the network architecture are illustrated schematically in Figure 6.4.

Due to the complexity of these data, we do not train a neural density estimator directly on the density maps. We instead first apply a CNN to extract features from our images and embed the corresponding information in a lower-dimensional latent vector. Following a similar procedure as in Chapter 5 (see also Ronchi et al., 2021), we choose the following baseline architecture for our embedding network:

- 2D convolution layer with kernel size 3×3, 3 input channels, 32 output channels, stride 1, padding 1.

- 2D Max pooling layer with size 2×2, stride 2, no padding.

- 2D convolution layer with kernel size 3×3, 32 input channels, 64 output channels, stride 1, padding 1.

- 2D Max pooling layer with size 2×2, stride 2, no padding.

- Fully connected linear layer with the flattened output from the second pooling layer as input and 32 output neurons encoding the latent representation.

After each convolution and the fully connected layer, we apply a rectified linear unit (ReLU) activation function (see Section 2.2.3). The weights for the CNN are initialised using the Kaiming prescription (He et al., 2015) to avoid exploding or vanishing gradients during the training process (see Section 2.3.2).

We subsequently pass the latent vector generated by the CNN to a neural density estimator. We implement a MDN and specifically opt for a Gaussian-mixture model (GMM) in five dimensions to approximate the posterior, $q_{F(x,w)}(\theta)$, for our five free magneto-rotational parameters. This implies

$$q_{F(x,w)}(\theta) = \sum_{c=1}^{C} \alpha_c \, \mathcal{N}(\theta | \mu_c, \Sigma_c), \tag{6.27}$$

where, C denotes the total number of Gaussian components used, α_c is the mixture weight and $\mathcal{N}(\theta | \mu_c, \Sigma_c)$ the multi-variate Gaussian distribution with mean vector μ_c and covariance matrix Σ_c for the c-th component.

For our MDN, we follow sbi's default implementation and use:

- Three fully connected layers with 32 neurons each.

- Four fully connected output layers which encode the Gaussian mixture weights, α_c, means, μ_c, diagonal and upper triangular components of the covariance matrices, Σ_c, respectively. These contain c, $5c$, $5c$ and $10c$ neurons, respectively.

We again apply the ReLU activation function after each hidden layer, while weights for the MDN are initialised with PyTorch's default initialisation (Glorot and Bengio, 2010).

We subsequently train the entire pipeline using the gradient descent optimiser Adam (Kingma and Ba, 2014, see also Section 2.3.4). At each epoch the network undergoes a series of optimisation steps based on the information provided in the entire training data set before epoch-averaged training and validation metrics are computed based on the negative of Equation (6.26). In this way during optimisation the algorithm tend to maximise the training and validation metrics. Note that we also set an early stop of 20 to prevent overfitting, which implies that the training process is interrupted (and the weights of the best validation epoch recorded) once the validation metric has not improved for 20 epochs.

TABLE 6.2: Information for the 22 machine-learning experiments summarising the specific training data and hyperparameters as well as the resulting metrics. The columns report the experiment number (#), the resolution (res) for our P-$\dot{P}$ density maps, the included surveys and the fraction (frac; in %) of the 291,600 populations in the training set used for training, information on whether we standardised (std) or normalised (norm) the input, the number of Gaussian components (comp) in our MDN, the batch size (BS), the learning rate (LR), the CNN architecture (we distinguish our baseline set-up and a deeper network; see Secs. 6.3.2 and 6.3.3 for details), the best metric computed over the validation set (VM), the number of training epochs, the time it took to train the network in seconds and the average metric computed over our 3,600 test samples (TM). In bold, we highlight those parameters that we have varied with respect to the baseline experiment #1. Experiments with an asterisk ($\star$) are removed from following analysis due to training irregularities.

#	res	surveys	frac (%)	input	comp	BS	LR	CNN	VM	epochs	time (s)	TM
1	32	PMPS, SMPS, HTRU	100	std	10	8	0.0005	baseline	3.65	38	9,373	3.64
2	32	PMPS, SMPS, HTRU	100	std	10	8	0.0005	**deep**	3.71	49	14,292	3.71
3	**64**	PMPS, SMPS, HTRU	100	std	10	8	0.0005	baseline	3.55	55	78,837	3.54
4	**64**	PMPS, SMPS, HTRU	100	std	10	8	0.0005	**deep**	3.64	89	128,119	3.64
5	32	PMPS, SMPS, HTRU	**75**	std	10	8	0.0005	baseline	3.74	71	13,232	3.78
6	32	PMPS, SMPS, HTRU	**50**	std	10	8	0.0005	baseline	3.56	58	7,000	3.55
7$\star$	32	PMPS, SMPS, HTRU	100	**norm**	10	8	0.0005	baseline	3.47	30	7,445	3.73
8	32	PMPS, SMPS, HTRU	100	**norm**	10	8	**0.01**	baseline	9.66	54	13,015	9.60
9	32	PMPS, SMPS, HTRU	100	std	**8**	8	**0.001**	baseline	3.74	52	12,389	3.73
10	32	PMPS, SMPS, HTRU	100	std	**5**	8	0.0005	baseline	3.83	118	27,973	3.86
11	32	PMPS, SMPS, HTRU	100	std	10	**16**	0.0005	baseline	3.99	85	10,476	3.97
12	32	PMPS, SMPS, HTRU	100	std	10	**32**	0.0005	baseline	4.11	79	5,346	4.06
13	32	PMPS, SMPS, HTRU	100	std	10	8	**0.001**	baseline	3.36	61	14,785	3.33
14	32	PMPS, SMPS, HTRU	100	std	10	8	**0.0001**	baseline	4.22	75	18,369	4.22
15	32	**HTRU**	100	std	10	8	0.0005	baseline	3.43	63	15,568	3.42
16	32	**SMPS, HTRU**	100	std	10	8	0.0005	baseline	3.58	40	9,979	3.59
17	32	**PMPS, SMPS**	100	std	10	8	0.0005	baseline	3.41	69	16,937	3.41
18$\star$	**64**	PMPS, SMPS, HTRU	**50**	std	10	8	0.0005	baseline	3.45	47	5,766	3.44
19	32	PMPS, SMPS, HTRU	100	**norm**	10	**32**	**0.001**	baseline	10.05	44	2,864	10.20
20	32	PMPS, SMPS, HTRU	100	**norm**	10	**32**	**0.0001**	baseline	10.31	90	5,815	10.49
21	32	PMPS, SMPS, HTRU	100	**norm**	10	**16**	**0.001**	baseline	9.82	77	9,901	9.98
22$\star$	32	PMPS, SMPS, HTRU	100	**norm**	10	**16**	**0.0001**	baseline	10.45	124	15,603	10.55

6.3.3 Experiments

Table 6.2 summarises the 22 different experiments that we have conducted for this study to assess the performance of SBI for pulsar population synthesis. For this purpose, we varied aspects of the training data as well as the hyperparameters of our deep-learning pipeline. In particular, for the input we explored two different resolutions for the P-$\dot{P}$ maps, 32 and 64, respectively, assessed the network performance when all three density maps or only two/one are provided, and whether normalisation or standardisation during preprocessing leads to different results. We further studied the impact of using the full training data set or smaller subsets. Moreover, for the network we varied the number of Gaussian mixture components in our neural density estimator, the batch size, and the learning rate, and we explored two different CNNs for our embedding net. In addition to the baseline architecture described in Section 6.3.2, we also conducted two experiments with a deeper network composed of four convolutional blocks. Here, the two convolutional layers introduced previously are followed by an additional layer with 32 and 64 input/output channels, respectively. Kernel size, stride, padding, subsequent pooling and fully-connected layers were kept as above.

Due to the computational cost of each training experiment, a full grid search over all relevant configurations was beyond the scope of this work. We, therefore, opted to produce a representative set of experiments that provide sufficient information to study the variation of our inferred posteriors in Section 6.4. Finally note, that almost all of our optimisations are performed on a Tesla V100 SXM2 GPU with 32 GB memory. We only trained experiments #3 and #4, for which the full training data set with a resolution of 64 was too large to be optimised on the GPU, on a CPU with 32 GB RAM. In those two cases, training the network, thus, took markedly longer than for the other experiments (see below).

6.4 Results

6.4.1 Training

Several metrics for our experiments are summarised in the last four columns of Tab. 6.2. We observe that the optimisation of our neural networks take between $\sim 1 - 8\,$hr on the GPU and on the order of a day on a CPU, completing $\sim 30 - 124$ training epochs. In general, we find good training behaviour with the validation metric closely tracking the training metric and little or no overfitting. This is also evident in the network's generalisation ability illustrated by the average metrics computed over the unseen test set of $3,600$ simulations. The evolution of the training and validation metrics for experiment #1 is shown in Fig. 6.5 as an example. We remind the reader that we aim to maximise the total log-posterior. After visual inspection of all training curves, we remove experiment #7 due to irregularities in the training behaviour and experiments #18 and #22 due to a slight tendency to overfitting. Note that these shortcomings were not directly visible from the training metrics in Table 6.2. We also highlight that we find systematically larger training, validation and test metrics in those experiments where our input density maps were normalised. In the following, we however assess the quality of the corresponding posteriors and find that these do not result in better inferences. Beyond this difference, we cannot identify any significant variation in the metrics between the remaining configurations. We, hence, proceed with an analysis of all experiments apart from numbers #7, #18 and #22.
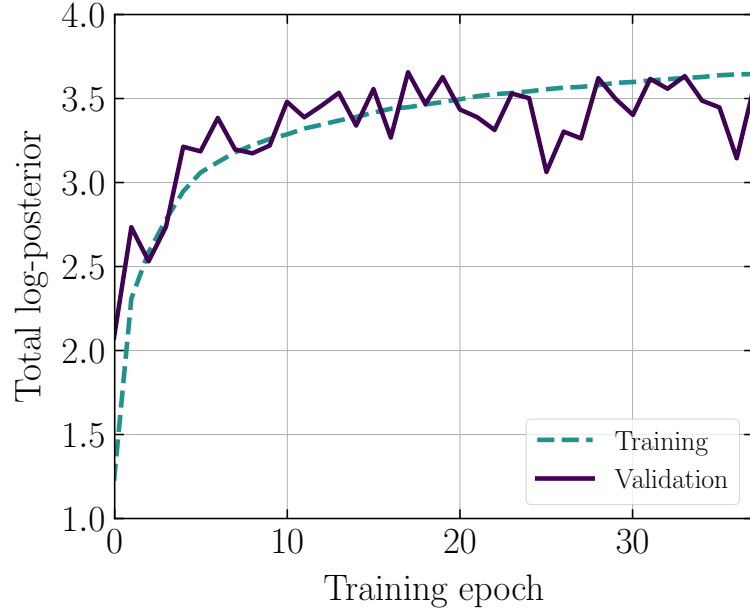
FIGURE 6.5: Training behaviour for baseline experiment #1. We show training metric (dashed, light blue) and validation metric (solid, purple) as a function of the training epoch. We seek to maximise the total log-posterior, $\sum_{i=1}^{N} \log q_{F(x_i, w)}(\theta_i)$, over the training and validation data set, respectively, as the network learns. Both losses increase as expected and the validation curve closely tracks the training curve, i.e., we see little overfitting. The best validation loss is reached at epoch 17 and the early stop criterion, thus, causes the training halt after 37 training epochs.

TABLE 6.3: Magneto-rotational parameters for three random test samples and the observed pulsar population. The first five rows show the ground truths, θ, used to simulate the test populations. The second block gives the median and 95% CIs obtained from inferences with the neural network trained in experiment #1. The final block contains the median and 95% CIs determined from the ensemble posterior combining 19 experiments.

| **Parameters**		**Test sample** 1	**Test sample** 2	**Test sample** 3	**Observed population**
Ground truths, θ	$\mu_{\log B}$	13.19	13.86	13.35	-
	$\sigma_{\log B}$	0.96	0.88	0.24	-
	$\mu_{\log P}$	-0.85	-0.42	-1.25	-
	$\sigma_{\log P}$	0.51	0.61	0.60	-
	a_{late}	-0.86	-1.71	-2.38	-
95% CI experiment #1	$\mu_{\log B}$	$13.28^{+0.18}_{-0.18}$	$13.73^{+0.15}_{-0.15}$	$13.33^{+0.05}_{-0.04}$	$13.07^{+0.07}_{-0.08}$
	$\sigma_{\log B}$	$0.95^{+0.08}_{-0.08}$	$0.79^{+0.07}_{-0.07}$	$0.23^{+0.02}_{-0.02}$	$0.43^{+0.03}_{-0.03}$
	$\mu_{\log P}$	$-0.90^{+0.13}_{-0.13}$	$-0.35^{+0.19}_{-0.18}$	$-1.17^{+0.33}_{-0.34}$	$-0.98^{+0.25}_{-0.29}$
	$\sigma_{\log P}$	$0.49^{+0.10}_{-0.09}$	$0.73^{+0.20}_{-0.15}$	$0.73^{+0.25}_{-0.31}$	$0.54^{+0.33}_{-0.25}$
	a_{late}	$-0.83^{+0.06}_{-0.06}$	$-1.88^{+0.35}_{-0.35}$	$-2.47^{+0.43}_{-0.43}$	$-1.77^{+0.35}_{-0.38}$
95% CI ensemble	$\mu_{\log B}$	$13.29^{+0.20}_{-0.20}$	$13.74^{+0.19}_{-0.16}$	$13.34^{+0.05}_{-0.05}$	$13.10^{+0.08}_{-0.10}$
	$\sigma_{\log B}$	$0.96^{+0.07}_{-0.08}$	$0.78^{+0.09}_{-0.08}$	$0.24^{+0.02}_{-0.02}$	$0.45^{+0.05}_{-0.05}$
	$\mu_{\log P}$	$-0.92^{+0.16}_{-0.15}$	$-0.40^{+0.20}_{-0.27}$	$-1.23^{+0.33}_{-0.34}$	$-1.00^{+0.26}_{-0.21}$
	$\sigma_{\log P}$	$0.49^{+0.10}_{-0.09}$	$0.74^{+0.20}_{-0.17}$	$0.67^{+0.30}_{-0.28}$	$0.38^{+0.33}_{-0.18}$
	a_{late}	$-0.84^{+0.06}_{-0.07}$	$-1.76^{+0.39}_{-0.43}$	$-2.34^{+0.43}_{-0.45}$	$-1.80^{+0.65}_{-0.61}$
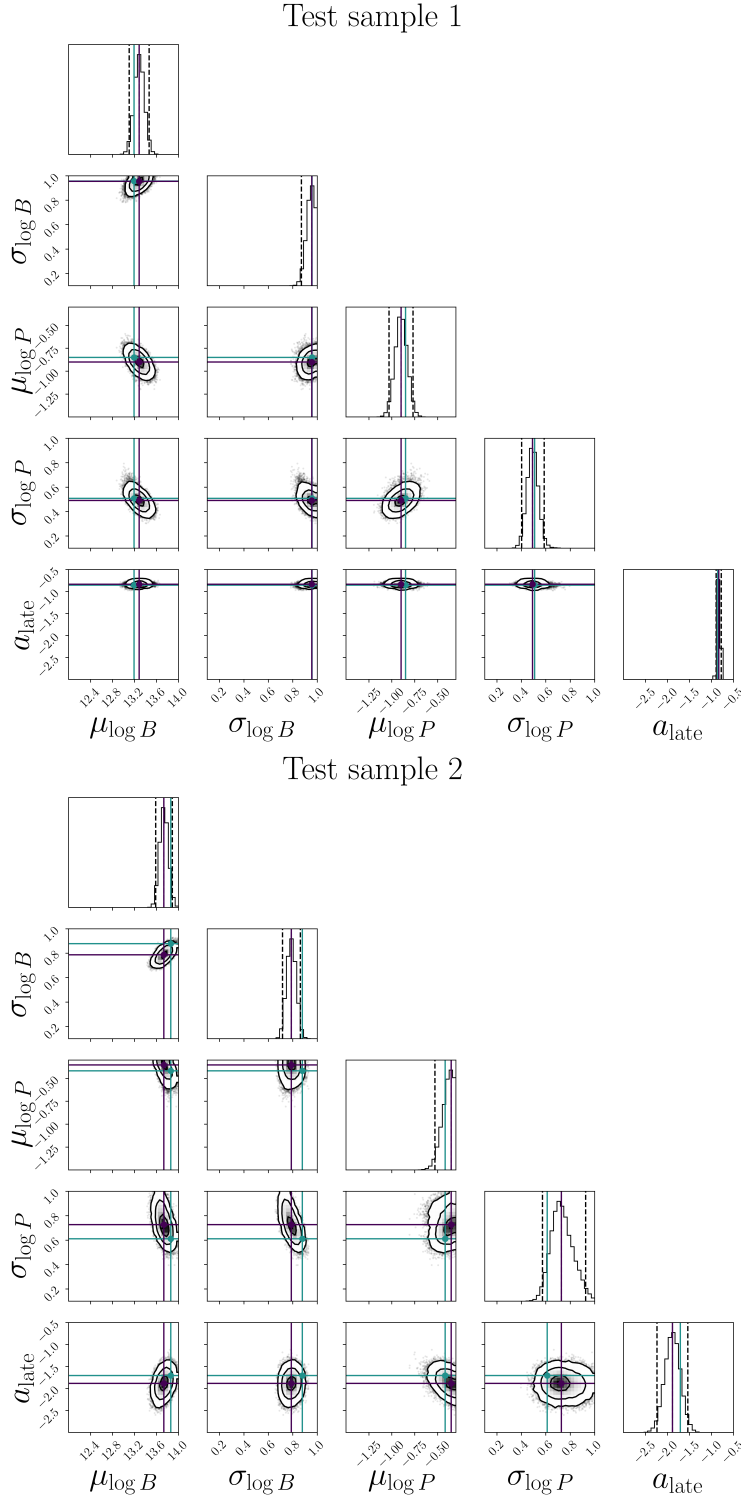
FIGURE 6.6: Benchmark inference for test simulation 1 and 2 using the network from experiment #1. The corner plot shows one- and two-dimensional marginal posterior distributions for the five magneto-rotational parameters. We also show corresponding ground truths, θ, in light blue and the medians in purple. We observe that the posteriors cover the θ well. Corresponding 95% CIs are summarised in Table 6.3.
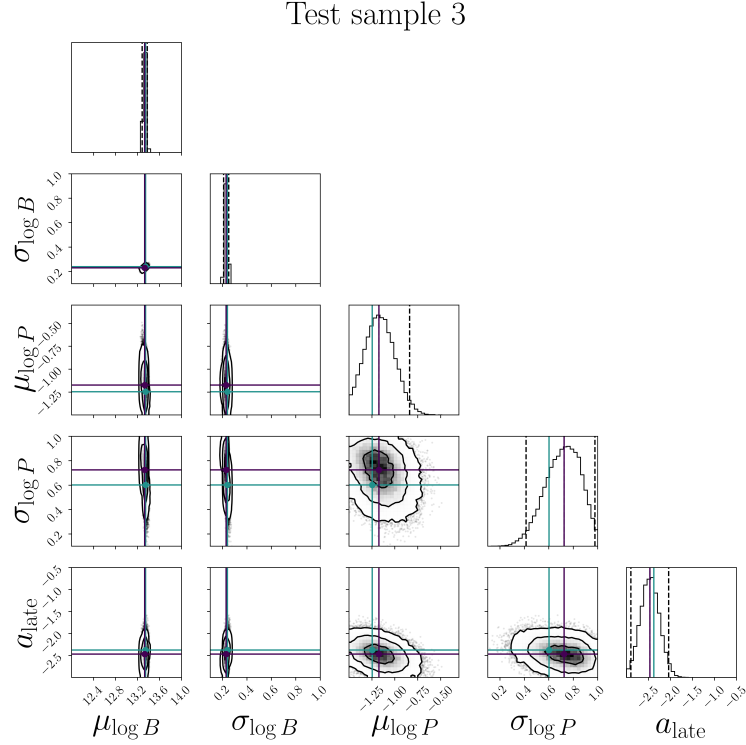
FIGURE 6.7: Same as Fig. 6.6 but for test simulations 3.

6.4.2 Benchmark inferences

As a first assessment of our approximated posteriors, we focus on inferring the five magneto-rotational parameters, $\mu_{\log B}, \sigma_{\log B}, \mu_{\log P}, \sigma_{\log P}, a_{\text{late}}$, for simulated populations where we know the input parameters, $\boldsymbol{\theta}$. We specifically look at the three simulations, whose P-$\dot{P}$ diagrams were illustrated in the top row of Figure 6.3. Corresponding ground truths, $\boldsymbol{\theta}$, are summarised in the top five rows in Table 6.3. In Figures 6.6 and 6.7, we show the resulting one- and two-dimensional marginal posterior distributions obtained by repeatedly sampling from the neural network optimised during experiment #1. For all three cases, the posteriors are well defined, significantly smaller than our prior ranges (6.24) shown along the axes, and centred around the ground truths, $\boldsymbol{\theta}$, highlighted in light blue. To quantify this, we calculate the 1σ, 2σ and 3σ credible regions, shown as contours in the two-dimensional posteriors. In the one-dimensional posterior panels, the corresponding 95% credible intervals (CIs) are given as dashed, black lines, while medians are illustrated as solid, purple lines. Their numerical values are given in Table 6.3. We observe that the ground truths, $\boldsymbol{\theta}$, are typically contained within the 2σ credible regions, which we interpret as evidence that our NPE approach is capable of producing reasonable posterior distributions. In general, the credible regions for the two parameters characterizing the initial magnetic-field distribution are narrower than those for the initial period distribution and the late-time magnetic-field decay. We confirm that the behaviour is qualitatively similar for the remaining P-$\dot{P}$ simulations in our test set.

We next compare the inferences for our various training experiments. To visualise corresponding differences, we plot the one-dimensional marginalised posteriors for all 19 experiments for the three test samples in grey in Figure 6.8. Ground
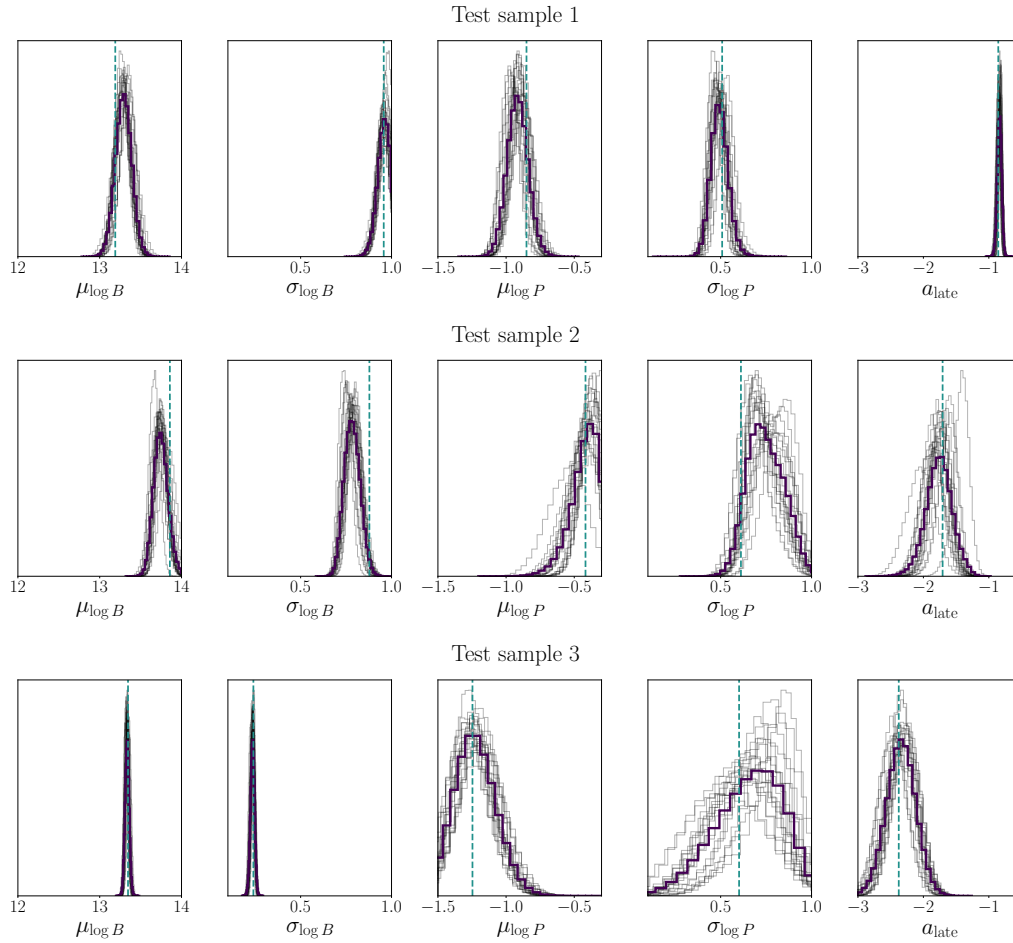
FIGURE 6.8: One-dimensional marginal posteriors for the five magneto-rotational parameters for the three test simulations inferred using 19 different NPE experiments shown in grey. The horizontal axes represent the parameters' prior ranges. The ground truths are shown as vertical dashed lines in light blue. We observe variation between the experiments, specifically for $\mu_{\log P}$, $\sigma_{\log P}$ and a_{late}. We also plot the ensemble posteriors (purple) obtained as a weighted average of the individual posteriors.

truths, $\boldsymbol{\theta}$, are shown as dashed lines in light blue. We observe that the width of individual posterior approximations as well as their medians can vary somewhat between different test samples and magneto-rotational parameters. Compared across the full test set, this behaviour is again more dominant for the period and late-time magnetic-field parameters than for the initial *B*-field properties. However, no individual NPEs stand out by exhibiting either particularly good or poor posteriors. Further note that we also do not see any differences for those experiments with normalised input maps that showed systematically better metrics than those experiments trained on standardised data. This highlights that training behaviour alone does not provide sufficient information on the quality of the resulting inference.

In light of this, we also determine the combined posterior for all 19 experiments (Hermans et al., 2021). We calculate the corresponding *ensemble posterior*, $\overline{q}(\boldsymbol{\theta})$, as the weighted average of the individual posteriors:

$$\overline{q}(\boldsymbol{\theta}) = \sum_j^{19} w_j \log q_{F_j}(\boldsymbol{\theta}), \tag{6.28}$$

where w_j represents the weight of the *j*-th component. Giving equal importance to each experiment in the ensemble, we choose $w_j = 1/19$. The corresponding one-dimensional marginalised ensemble posteriors for $\mu_{\log B}$, $\sigma_{\log B}$, $\mu_{\log P}$, $\sigma_{\log P}$ and a_{late} for the three test simulations are illustrated as purple histograms in Figure 6.8. As expected, they fall within the individual posteriors. The corresponding 95% CIs for the three test samples, which are typically comparable or slightly wider than those calculated for experiment #1 posteriors alone, are summarised in the bottom five rows of Table 6.3.

6.4.3 Posterior validation

To further assess whether posterior estimates are well-calibrated, we determine their *coverage* (see Section 2.5.2). As outlined in detail in Appendix D, the coverage probability measures the fraction of test samples for which (for a given credibility level $1 - \alpha$) the ground truths, $\boldsymbol{\theta}$, fall within the corresponding $1 - \alpha$ region of their respective posteriors, $q_{F(\boldsymbol{x}, \boldsymbol{\phi})}(\boldsymbol{\theta})$. For a well-calibrated posterior distribution and a sufficiently large number of test samples, this fraction should equal $1 - \alpha$. This implies that the coverage probability as a function of the credibility level is diagonal. In contrast, for a conservative posterior that is wider than the true posterior, we would recover a fraction larger than $1 - \alpha$. Conversely, for a narrower (overconfident) posterior, the corresponding fraction of test samples is less than $1 - \alpha$. In terms of the coverage, this corresponds to curves above and below the diagonal, respectively, and can, therefore, be used to assess the quality of approximate posteriors.

We show the coverage probabilities for our different posterior estimates as a function of the credibility level, $1 - \alpha$, in Figure 6.9. We specifically highlight the coverage for the posterior from experiment #1 (dashed, light blue) and the ensemble posterior (solid, purple). All remaining experiments are shown in grey. We observe that the approximate posteriors for individual experiments closely follow the diagonal, exhibiting either slightly conservative or overconfident behaviour. As expected, the most conservative estimate is given by our ensemble posterior, which incorporates variations in the inference for 19 different machine-learning configurations across all $3{,}600$ test samples. These results provide additional support that our neural posteriors are trustworthy and have indeed learned to accurately infer magneto-rotational parameters from simulated P-$\dot{P}$ density maps.
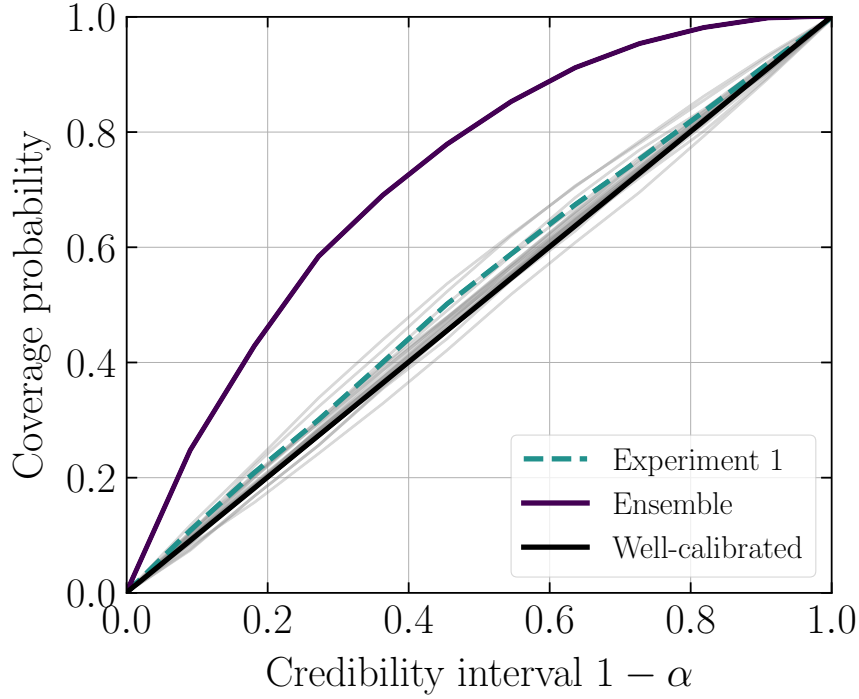
FIGURE 6.9: Coverage probability as a function of the credibility level, $1 - \alpha$, for our approximate posteriors calculated for the $3,600$ test simulations. We specifically highlight the coverage for experiment #1 as a dashed, light blue line and that for the ensemble as a solid, purple line. All remaining experiments are given in grey. For a well-calibrated posterior, the coverage follows the diagonal shown in black.
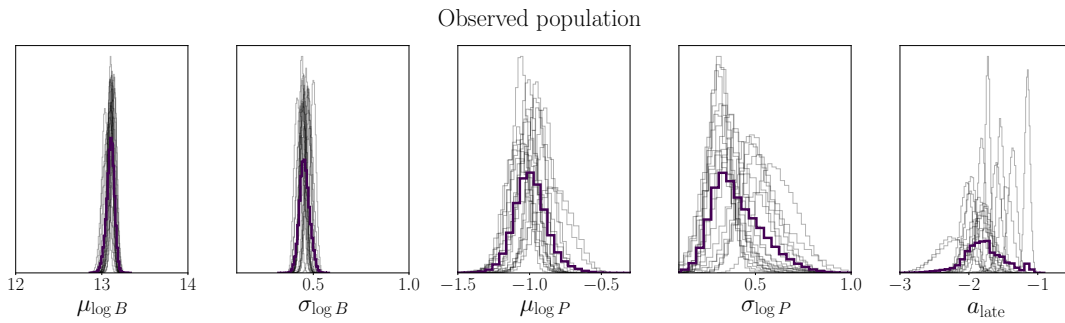

FIGURE 6.10: One-dimensional marginal posteriors for the five magneto-rotational parameters for the observed pulsar population. We show inference results for 19 different NPE experiments in grey and the ensemble posterior in purple.

6.4.4   Inference on the observed population

Following the benchmark experiments and the coverage determination, we now turn our attention to the true pulsar populations observed with the PMPS, the SMPS and the low- and mid-latitude HTRU survey. The corresponding P-$\dot{P}$ diagram was shown in the bottom panel of Fig. 6.2. We represent these populations as three density maps as outlined in Section 6.2.5 and subsequently feed them through our trained neural networks to infer the five parameters, $\mu_{\log B}, \sigma_{\log B}, \mu_{\log P}, \sigma_{\log P}$ and a_{late}, assuming that our simulation framework provides a realistic description of the underlying physics.

We show the corresponding one-dimensional marginal posterior distributions for individual experiments (grey histograms) and the ensemble (purple histograms) in Figure 6.10. Additionally, a corner plot for the one- and two-dimensional ensemble posteriors is illustrated in Figure 6.11. Corresponding medians (shown in purple in the corner plot) and 95% CIs for experiment #1 and the ensemble are also summarised in the last column of Table 6.3.

The general trend (already observed for the simulated populations) that the initial magnetic-field parameters, $\mu_{\log B}$ and $\sigma_{\log B}$, are much better constrained by our NPE framework than the remaining three values also holds for the observed population. As seen in the first two panels of Figure 6.10, all 19 experiments recover narrow posteriors around similar medians. For the initial period-distribution parameters, $\mu_{\log P}$ and $\sigma_{\log P}$, (see third and fourth panel, respectively), we obtain wider posteriors and a larger variety of median values between different experiments. These posteriors, however, cover similar regions within our prior ranges and are comparable to what we observed for the test samples. In contrast, the inferred posteriors for a_{late} (the final panel in Figure 6.10) exhibit different behaviour to our benchmark experiments. In particular, posteriors vary significantly in width between different experiments with those at the larger (smaller) end of the a_{late} range generally exhibiting narrower (larger) widths. Moreover, several distributions do not overlap at all. This is manifest as a relatively wide posterior in the ensemble which also shows a second peak, primarily driven by the right-most individual posterior resulting from experiment #2. Note that this configuration did not cause irregularities during the network optimisation or unusual posteriors for our test samples. We, therefore, do not associate this behaviour with the network itself. The corresponding bi-modality is also visible in the final row of the corner plot in Figure 6.11. We will discuss our interpretation of this below.

6.5   Discussion and conclusions

In this study, we have successfully developed a new machine-learning pipeline that combines pulsar population synthesis with simulation-based inference (SBI) for the first time and tested the corresponding approach for inferring magneto-rotational properties of neutron stars.

6.5.1   Simulation framework

We first discussed our implementation of the forward model, i.e., the prescription for simulating the dynamical and magneto-rotational properties of the Galactic population of isolated radio pulsars, modelling their radio emission and subsequently mimicking observational limitations for the Parkes Multibeam Pulsar Survey (PMPS),
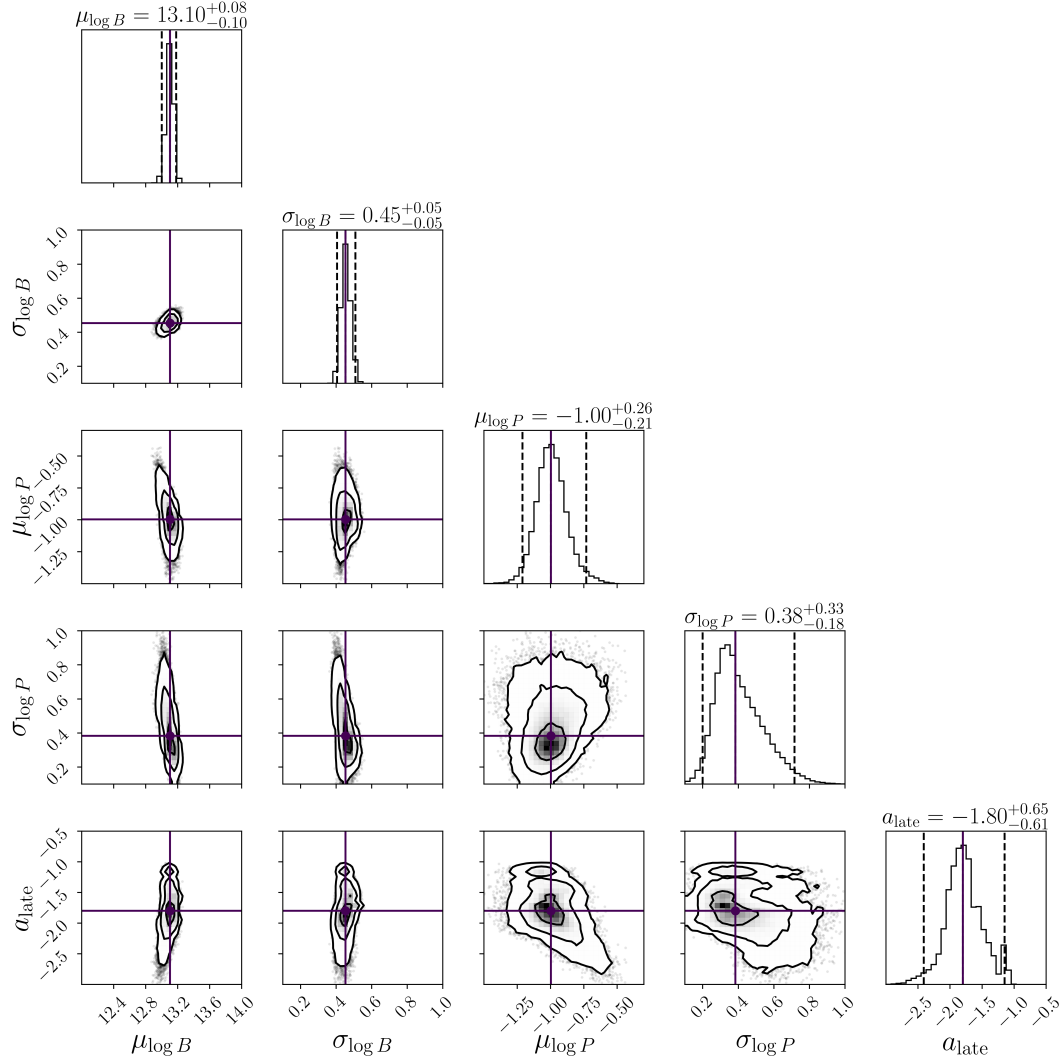
FIGURE 6.11: Inference results for the observed pulsar population using the ensemble posterior of 19 different NPEs. The corner plot shows one- and two-dimensional marginal posterior distributions for the five magneto-rotational parameters. We highlight the medians in purple. Corresponding values and 95% CIs are summarised above the panels and in Table 6.3.

the Swinburne Intermediate-latitude Pulsar Survey (SMPS) and the low- and mid-latitude High Time Resolution Universe (HTRU) survey. We followed earlier frameworks (e.g., Faucher-Giguère and Kaspi, 2006; Bates et al., 2014; Gullón et al., 2014; Gullón et al., 2015; Cieślar, Bulik, and Osłowski, 2020) but implemented several key differences (see also Table 1.1). In particular, we sampled the birth positions of our pulsars from the Galactic electron distribution (Yao, Manchester, and Wang, 2017) instead of following the typical approach of combining a spiral-arm model with a radial pulsar distribution of, e.g., Yusifov and Küçük (2004). The latter is deduced for the observed, evolved pulsar sample and not the initial population. Moreover, we have included the (rigid) rotation of the Galaxy, which leads to a more consistent treatment of the pulsar birth positions compared to earlier analyses (see also Chapter 5 and Ronchi et al., 2021). For the magnetic-field evolution, we used a similar approach to Gullón et al. (2014) and Gullón et al. (2015) taking advantage of the newest two-dimensional magneto-thermal simulations (Viganò et al., 2021) and solved for the coupled evolution of the spin period, P, and the misalignment angle, χ, for a plasma-filled magnetosphere. To capture the field changes at late times, we developed a new physically motivated prescription in which the magnetic field, B, decays according to a power law captured by the index, a_{late}. Together with the means, $\mu_{\log B}, \mu_{\log P}$, and standard deviations, $\sigma_{\log B}, \sigma_{\log P}$, which characterise the normally distributed logarithms of the initial periods and the initial fields, we hence obtained five parameters that control the neutron stars' magneto-rotational evolution.

To simulate the detection of our synthetic pulsars, we make two main changes compared to earlier studies. First, we do not model the pulsars' pseudo luminosity (see Section 1.11) defined as $L_{f,\text{pseudo}} \propto S_{f,\text{mean}} d^2$ (where $S_{f,\text{mean}}$ is the period-averaged detected flux at frequency f, and d the pulsar distance) but instead assume that the intrinsic neutron-star luminosity, L_{int}, is proportional to the spin-down power, $\dot{E}_{\text{rot}}$ (see Section 1.9.4). In particular, we considered $L_{\text{int}} \propto \dot{E}_{\text{rot}}^{1/2}$ to determine the bolometric radio flux and subsequently propagate the corresponding pulsed emission towards the Earth. We also used a geometry-based description to determine the pulsars that are beamed towards us, which earlier works typically treat in an empirical manner. Finally, we not only looked at PMPS and SMPS but also incorporated the HTRU survey for the first time. Using the resulting simulation framework, we then produced $360,000$ synthetic P-$\dot{P}$ diagrams which we converted to one density map per survey in preparation for the neural network. 90% of these simulations were used for training and validation, and the remaining 10% reserved for testing.

6.5.2   Inference procedure

The second part of this study is centred on the implementation of the SBI approach, specifically focusing on neural posterior estimation (NPE), to learn a probabilistic association between our simulator output and the input parameters, $\boldsymbol{\theta} = \{\mu_{\log B}, \sigma_{\log B}, \mu_{\log P}, \sigma_{\log P}, a_{\text{late}}\}$. To do so, we first used a convolutional neural network (CNN) to extract features from our high-dimensional P-$\dot{P}$ maps and obtain a compressed representation, which was then transferred into a flexible neural density estimator. By taking advantage of the open-source Python package sbi (Tejero-Cantero et al., 2020),[4] we specifically implemented a Gaussian-mixture density model in five dimensions to approximate our posterior. To study the sensitivity of the NPE results

---

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

on the representation of our input data and the network hyperparameters, we conducted 22 distinct experiments. An inspection of the corresponding training metrics led us to discard three experiments due to irregular training behaviour or overfitting. The remaining 19 trained neural networks were analyzed further and we found no significant differences in the resulting inferences when benchmarked on three random test simulations. The same was observed when validating the posteriors through a coverage calculation over the test set with $3,600$ samples, highlighting that all 19 posterior estimates are well-calibrated. From this we concluded, in particular, that the training behaviour is a poor identifier of subsequent inference quality, because normalisation of input maps led to systematically better training, test and validation metrics compared to standardizing the input but comparable inferences. Learning rate and batch size played a negligible role in both set-ups.

We also point out that the use of smaller training data sets did not affect the inference quality either. While we expect that training sets of $\lesssim 10\%$ (i.e., $30,000$ simulations) will eventually have an effect on this, databases of 50% (i.e., $150,000$ simulations) are sufficient when inferring five parameters. For comparable studies, this would imply a significant reduction in simulation time, the most costly part of these analyses. Similar performances further justify optimizing our networks for density maps with a resolution of 32×32 bins instead of 64×64 and the shallower baseline CNN to speed up the training process. Additionally, we highlight that the use of different numbers of Gaussian mixture components also led to comparable optimisation metrics and inference results. Extracting the corresponding mixture weights, α_c, after the optimisation, we find that we only require two or three Gaussians to approximate our posteriors across the entire test data set. We do, however, point out that training with a larger number of components was faster. Finally, note that the use of fewer surveys (i.e., one or two density maps only) did not change the inference results for our five magneto-rotational parameters. Naively, one might think that complementary information on the pulsar population as, e.g., provided by SMPS, which is sensitive to older stars at higher Galactic latitudes, would help the network learn better posteriors. We do, however, not observe such behaviour in our experiments. Although this might suggest that using single surveys in the future could be sufficient to constrain neutron-star parameters through population synthesis, we caution that different surveys, in principle, provide additional information on the neutron-star birth rate (see below) which was not supplied to our neural networks, i.e., we focused on the location and shape of the pulsar population in the P-$\dot{P}$ plane only.

Due to the variations in our inference results, and because we could not identify a single neural network as the best posterior estimator, we also determined the ensemble posterior through an equally weighted average of the individual experiments. The resulting posterior behaved as expected and showed more conservative behaviour than the ensemble members. For the next section, we, will, hence, follow the recommendation by Hermans et al. (2021) and use our (most conservative) ensemble posterior to analyze the observed pulsar population.

6.5.3   Inference results on the observed population

Following the validation of our NPE approach, we subsequently used the ensemble posterior estimator to infer the five magneto-rotational parameters for the true population of isolated Galactic radio pulsars observed with our three surveys. In

particular, we found the following best estimates at 95% credible level:

$$\mu_{\log B} = 13.10^{+0.08}_{-0.10},$$
$$\sigma_{\log B} = 0.45^{+0.05}_{-0.05},$$
$$\mu_{\log P} = -1.00^{+0.26}_{-0.21}, \qquad\qquad (6.29)$$
$$\sigma_{\log P} = 0.38^{+0.33}_{-0.18},$$
$$a_{\mathrm{late}} = -1.80^{+0.65}_{-0.61}.$$

The corresponding corner plot was shown in Figure 6.11.

As noted during the benchmarking experiments, we generally obtain narrower posterior distributions for the initial magnetic-field parameters when compared to the initial period parameters. Difficulties in constraining rotational birth properties are, however, not a shortcoming of our inference approach itself as this was also noted by earlier population-synthesis analyses (e.g., Gullón et al., 2014; Gullón et al., 2015; Cieślar, Bulik, and Osłowski, 2020). Instead, this has a physical reason that lies in the stars' coupled magneto-rotational evolution. First of all, as seen in Section 3.2, the spin period evolution tends to loses memory of the initial spin periods for $t \gg t_{\mathrm{em}}$ where $t_{\mathrm{em}} \propto P_0^2 / B_0^2$. Moreover, while the B-field initially stays constant (see Figure 6.1), pulsars move from the top left in the P-$\dot{P}$ plane diagonally towards the bottom right. As soon as the field starts to decay, they divert from this diagonal and begin to evolve almost vertically towards smaller $\dot{P}$ values in a narrow range of periods. This effect is more pronounced for stronger initial B-fields and steeper late-time field decay and particularly visible for test simulation 3, whose P-$\dot{P}$ diagram was shown in the top right panel of Figure 6.3. This process essentially leads to the loss of information on the initial period and makes it harder to infer corresponding parameters. As expected, test sample 3, therefore, shows the largest 95% CIs for $\mu_{\log P}$ and $\sigma_{\log P}$ out of our three test samples (see third column in Table 6.3 and last row in Figure 6.8).

Contrasting the posterior medians from Equation (6.29) to other population-synthesis studies, we first note that our $\mu_{\log B}$ estimate is roughly consistent with Gullón et al. (2014) and Gullón et al. (2015) but somewhat larger than those of Faucher-Giguère and Kaspi (2006) and Cieślar, Bulik, and Osłowski (2020). Moreover, we obtain a smaller $\sigma_{\log B}$ than Gullón et al. (2014) and Gullón et al. (2015) and Faucher-Giguère and Kaspi (2006) but a slightly larger estimate than Cieślar, Bulik, and Osłowski (2020). Although all four studies determine optimal parameter ranges different to us, we expect the differences in the initial B constraints to be mainly due to our more realistic prescription for the magnetic-field and coupled P-χ evolution. A direct comparison between our initial period parameters and earlier literature values is not possible at this stage, because (following recent results by Igoshev et al. (2022); see also Xu et al. (2023)) we have considered the periods' logarithm and not the periods themselves to be normally distributed. We, however, highlight that our inferred initial period and magnetic-field parameters are comparable with those of Igoshev et al. (2022) obtained through a simplified analysis of 56 young neutron stars in supernova remnants. As this study looked at magneto-rotational properties only, the authors were able to define an explicit likelihood function and perform statistical inference. In this context, we also point out that although Cieślar, Bulik, and Osłowski (2020) derive (relatively narrow) posteriors for a range of pulsar properties using an MCMC analysis, their underlying simulation framework is significantly reduced compared to ours invoking, e.g., (unrealistic) exponential field decay, vacuum magnetospheres, no coupling between periods and misalignment angles, and
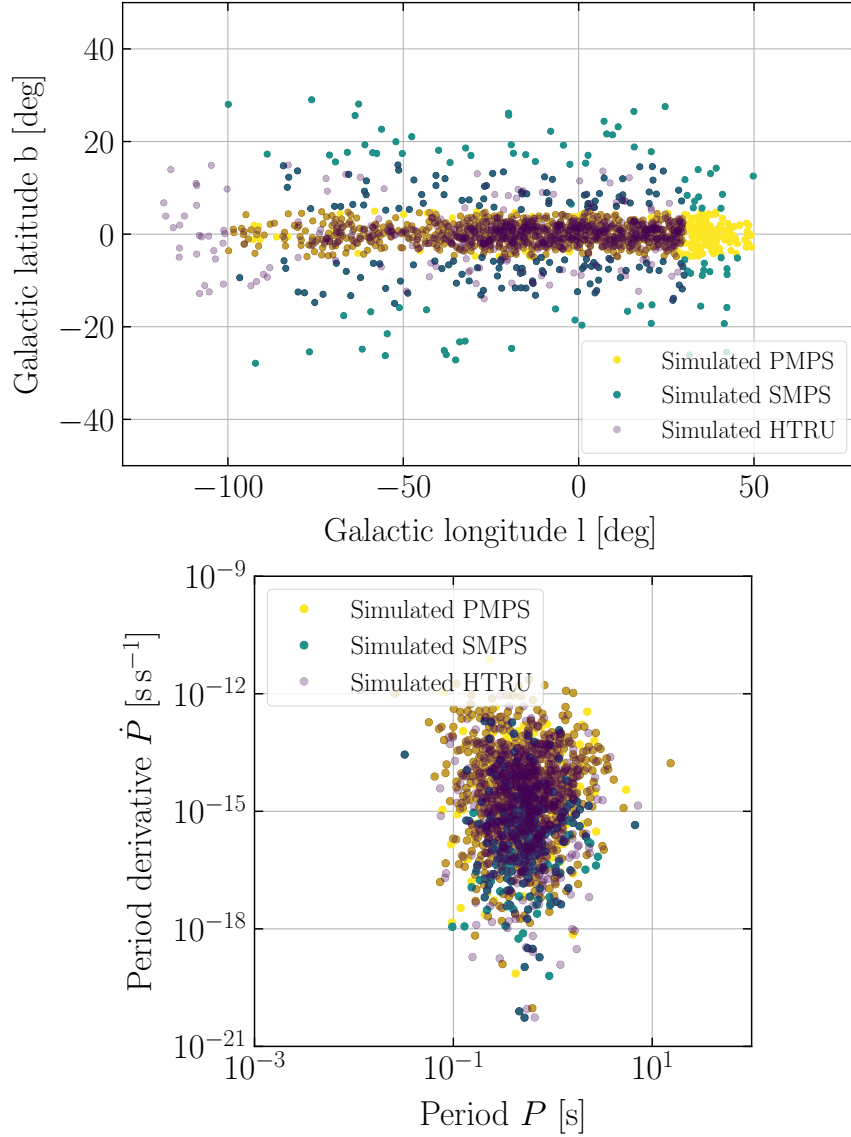
FIGURE 6.12: Simulated populations of isolated Galactic radio pulsars detected with the PMPS, the SMPS and the low- and mid-latitude HTRU survey (highlighted in yellow, light blue and purple, respectively) for the parameters inferred via SBI from the observed radio pulsar population (see Equation (6.29)). The left panel shows the distribution of the simulated population in Galactic latitude, b, and longitude, l, while the right panel depicts the pulsars in the period, P, and period derivative, $\dot{P}$, plane. Both plots directly compare to the true (observed) population shown in Figure 6.2.

a simplified prescription of the beamed emission. In addition, they make an explicit assumption on the likelihood that might not accurately capture the complexity of the pulsar population synthesis even for their simplified model. We reiterate the robustness of our SBI approach which eliminates the need for an explicit expression for the likelihood and is, therefore, also suitable for more complex simulators like ours. Moreover, as outlined above the use of a neural density estimator results in amortised posterior distributions that allow fast evaluation and sampling. We used this fact to determine the coverage and validate our posteriors, a procedure that is infeasible in MCMC approaches due to the time-consuming need for repeated sampling.

We now turn our attention to the parameter, a_{late}, the power-law index for the late-time magnetic-field decay. We newly introduced a_{late} in pulsar population synthesis to account for the highly uncertain, core-dominated field evolution above 10^6 yr in a phenomenological way. While corresponding inferences were satisfactory for our benchmark experiments, we found that posteriors for a_{late} inferred from the observed population differed significantly between our 19 experiments, systematically resulting in larger 95% CIs for smaller a_{late} medians and vice versa (see right most panel in Figure 6.10). In addition, several posteriors did not overlap at all across our prior range, leading to a bi-modality in the ensemble posterior. As we do not see anything similar for our synthetic simulations, we do not associate this behaviour with the networks' performance or the SBI approach itself. Instead, we hypothesise that this is due to shortcomings in our simulation framework. Put differently, our statistical inferences are only as good as the simulation model used to train our density estimator. Consequently, we see the complications in inferring a_{late} as an indication that our treatment of the late-time field evolution via a power law (albeit physically motivated by the behaviour of known magnetic-field evolution mechanisms) is insufficient to model the observed pulsar population.

Although further work is needed to better understand this aspect of neutron-star evolution, we can assure ourselves that our current power-law prescription is not too far away from the truth. To do so, we rerun our simulator with the best estimates summarised in Equation (6.29) and show the distribution of the resulting population in Galactic longitude and latitude and P and $\dot{P}$ in Figure 6.12. These are analogous to the respective plots in Figure 6.2. While a detailed comparison between this simulated and the observed population and a study of implications for the neutron-star birth rate is beyond the scope of this work, we will highlight a few main aspects. The distributions looked markedly similar, giving a reasonable level of confidence in our underlying simulation framework. However, we do see a slight shift in the SMPS population in the P-$\dot{P}$ diagram towards lower $\dot{P}$ values. This might again hint at missing physics at late times because SMPS is sensitive to somewhat older pulsars compared to the other two surveys. We further count the number of detected pulsars in the three synthetic surveys to obtain 1013, 239 and 1262 for the PMPS, the SMPS, and HTRU survey, respectively. Comparing these to the true observed counts in Equation (6.23), we find an equivalent number of objects in PMPS (within the sensitivity limits of our iterative approach of generating and detecting pulsars as summarised in Section 6.2.5), while we overestimate the SMPS population by $\sim 10\%$ and the HTRU population by $\sim 23\%$. To understand this discrepancy, we return to our earlier discussion of the neutron-star birth rate in Section 6.2.5. In particular, for our best-estimate simulation, we reach the observed

target counts (6.23) for each survey for the following birth rates:

$$\text{PMPS: } \sim 2.0 \text{ neutron stars per century,}$$
$$\text{SMPS: } \sim 1.9 \text{ neutron stars per century,} \qquad (6.30)$$
$$\text{HTRU: } \sim 1.7 \text{ neutron stars per century.}$$

Although these are similar, they are somewhat smaller than those obtained in earlier population-synthesis studies (Gullón et al., 2014; Faucher-Giguère and Kaspi, 2006) and very close to the recent core-collapse supernova estimate from Rozwadowska, Vissani, and Cappellaro (2021) (1.63 ± 0.46 per century). The difference in Equation (6.30) is sufficient to result in the overproduction of objects noted above. This is because we continue producing neutron stars until we hit the number of observed pulsars in all three surveys. In our specific case, PMPS detections require a slightly larger birth rate than the other two surveys. As mentioned previously, the main reason for this discrepancy is that we only expect the *correct* physical model to produce the same birth rate across all surveys. While we have successfully constrained magneto-rotational parameters for pulsar population synthesis using SBI for the first time, the birth-rate results (6.30) also suggest that our simulator has room for further improvement.

6.5.4 Future directions

In light of the previous conclusions, we intend to further develop our current approach in a number of ways. On the simulation side, we will investigate additional luminosity prescription that go beyond our assumption, $L_{int} \propto \dot{E}^{1/2}$, as this is another quantity that can significantly affect the pulsar distribution. Varying the exponent in our simulations, which was beyond the scope of this study due to computational limitations, and using SBI to constrain corresponding parameter ranges would be a first step in that direction. Additionally, we also assumed that the early-time *B*-field evolution is correctly captured by the two-dimensional magneto-thermal simulations from Viganò et al. (2021). Although the simulations we employed are based on reasonable assumptions of relevant physics and explain observations of a range of neutron stars, they are based on certain choices for the equation of state, crustal impurities, superfluid properties and magnetic-field configuration (see Appendix D for details). Analyzing the impact of these aspects on our statistical inference would be another useful step forward towards assessing the robustness of our results. As already noted above, further investigations into the field evolution of the neutron-star core at late times will also be crucial for future population-synthesis frameworks. Finally, new pulsar surveys (in the radio band as well as in other wavelengths) might hold the key to further constraining the neutron-star population. While we did not see a significant improvement in our inferences using information from one, two or three radio surveys, future studies will benefit from larger numbers of detected pulsars and accurate classification of telescope and detection biases. Furthermore, other wavebands, specifically gamma- or X-rays, provide complementary information that might help break degeneracies, albeit at the cost of introducing additional physics.

The increase in simulator complexity associated with these improvements will not only result in more free parameters but also inevitable lead to larger computation times for our forward model. The approach taken here, i.e., simulating a large database for input parameter combinations that cover the entire space sufficiently, will become infeasible. To overcome these hurdles, we will also have to explore new

SBI approaches. Sequential methods (e.g. Papamakarios, Sterratt, and Murray, 2018; Deistler, Goncalves, and Macke, 2022; Bhardwaj et al., 2023) that reduce the need for simulations by starting from a relatively small database and adaptively providing additional simulations (generated for those parts of the parameter space that are most useful for a neural density estimator to learn a posterior approximation) seem particularly suited to these tasks.

Chapter 7

Summary and conclusions

In this last chapter, I summarise the main results of this work, separately addressing the two main areas of this thesis, and outline the future directions.

7.1 Towards a unified evolutionary scenario for the neutron-star population

For this thesis we developed a flexible population-synthesis framework that is able to model the neutron-star birth properties and evolution in the Galaxy, starting from different prescriptions for the initial distributions in Galactic height, kick velocity, spin period, magnetic-field strength and magnetic-field evolution. Moreover, we simulated the neutron stars' electromagnetic emission in the radio band and modelled observational biases to emulate the detection with current radio facilities (see Chapters 1, 5 and 6). By varying the underlying model parameters, our simulation framework is able to generate different mock neutron-star populations that can be compared with the real observed population. To do this, we combine for the first time pulsar population synthesis with a machine learning framework with the aim of performing parameter inference on the observed data and constraining our physical models. Neural networks represent a powerful tool to perform this inference task due to their potential to process multi-dimensional data and extract relevant features in an automated way (see Chapter 2).

In Chapter 5 we focused on analysing the dynamics of Galactic neutron stars via simulations. In particular, by varying the natal kick-velocity distribution and the distribution of birth distances from the Galactic plane we evolved the pulsar trajectories in time, and generated a series of simulated populations. We then studied the feasibility of using convolutional neural networks to infer the dynamical properties at birth of a simulated population of neutron stars from density maps storing the information on their sky positions and proper motions. This proof-of-concept work showed that neural networks have the potential to predict the parameters governing the kick-velocity and Galactic height distributions with very good accuracy. The current limitation of this approach is the lack of observational data as we only know around 200 neutron stars with measured sky positions and proper motions. We demonstrated that a good inference accuracy would require about 10 times the current number of neutron stars with measured proper motions. Our work highlights the crucial need for detecting more neutron stars and accurately classifying them, a goal that can be reached with future radio telescopes such as the Square Kilometer Array. A future extension of this analysis should focus on the possibility of testing different kick-velocity distributions as recent studies have hinted at a possible bi-modality in the kick-velocity distribution due to different mechanisms producing the kicks in the supernova explosion.

In Chapter 6, we further improved our framework and focused on modelling the magneto-rotational evolution, the radio emission and the detection of neutron stars with current radio surveys. This study aimed to constrain the birth distributions of spin periods and magnetic fields and how these properties evolve in time for the radio pulsar population. We primarily focused on the radio pulsar population as we have a wealth of data from large-scale surveys performed with Murriyang (the Parkes radio telescope) in Australia in the past decades. As we have a fairly good understanding of the biases and systematics that affect the observed sample from these surveys, they represent a safe testing ground for our deep-learning approach.

We specifically explored an approach known as simulation-based inference or likelihood-free inference which represents a step forward from the approach explored in Chapter 5. There our analysis was limited to point estimates of dynamical neutron-star properties. With the approach presented in Chapter 6 we can instead evaluate a posterior distribution for the parameters. As introduced in Chapter 2, in simulation-based inference a deep neural network is employed as a probability density estimator that directly maps a given simulation outcome to the posterior probability distribution of the underlying model parameters without the need to compute the likelihood. This is particularly useful for complex simulators with multidimensional parameter inputs like ours, where defining and computing a likelihood probability distribution can be complicated or even intractable. To train the network for our specific problem, we employed density maps containing information on the spin period and spin period derivative of mock observed pulsar populations simulated with different distributions of initial magnetic fields and spin periods and different power laws for the magnetic-field evolution at late times. The trained neural network is then applied to the real empirical data to derive the posterior distribution of the parameters describing the real observed population of radio pulsars. We found that the observed radio pulsars can be best described by belonging to an underlying population born with the logarithm of the initial magnetic field following a normal distribution with best parameters $\mu_{\log B} = 13.10^{+0.08}_{-0.10}$, $\sigma_{\log B} = 0.45^{+0.05}_{-0.05}$ and the logarithm of initial spin period following a normal distribution with best parameters $\mu_{\log P} = -1.00^{+0.26}_{-0.21}$, $\sigma_{\log P} = 0.38^{+0.33}_{-0.18}$. We also put the first constraint on the magnetic-field evolution at late times. Our result suggests that assuming a late-time power law evolution, the magnetic field should decay with a power law-index $a_{\text{late}} = -1.80^{+0.65}_{-0.61}$.

A further benefit of the simulation-based inference approach is that the robustness of the results and their uncertainty can be checked using the coverage probability test (see Chapter 2 and Appendix D). This allows us to determine if the inferred posterior is well-calibrated or shows signs of over-confidence or under-confidence. By applying this test to our trained posterior density estimator we showed that the predicted posteriors are well-calibrated and our results using a posterior-ensemble are conservative.

The work completed for this thesis will form the basis to further develop a simulation framework able to include the high-energy emission of neutron stars. To this end, we plan to consider the results of new magneto-thermal 3D simulations (Gourgouliatos, Wood, and Hollerbach, 2016; De Grandis et al., 2020; Dehman et al., 2023) to explain the thermal X-ray radiation from the neutron-star surface and a synchro-curvature model (Viganò and Torres, 2015; Viganò, Torres, and Martín, 2015) to describe the high-energy gamma-ray radiation originating in the neutron-star magnetospheres. We also seek to emulate the observational biases and detection with current X-ray and gamma-ray facilities in order to compare our models with the

observed data. This multi-wavelength approach to studying the neutron-star population will allow us to consider and analyse the properties of different neutron star classes in a unified scenario. This will be particularly helpful to remove degeneracies and better constrain the parameters describing the neutron-star population as whole.

Furthermore, we will further improve our simulation-based inference framework by including an active learning approach. After training the density estimator over an initial set of simulations (much smaller than what we outlined in Chapter 6), the resulting initial estimate of the posterior distribution obtained for the observed data can be used to generate additional simulations that more closely reproduce the observed data. This new dataset of simulations can then be adopted as a new training set to refine the estimate of the posterior sequentially. This will allow a more effective inference procedure by concentrating the computational resources on the regions of the parameter space that is more compatible with the observed data.

7.2 The mystery of long-period pulsars

The second topic that I have explored in this thesis concerns the recent discovery of three intriguing sources: PSR J0901-4046, the slowest radio neutron-star pulsar ever detected with a spin period of 76 s and two mysterious periodic radio sources, GLEAM-X J1627 and GPM J1839–10, with periods of 1091 s and 1318 s, respectively. The emission properties of these sources are similar to those of magnetars but the nature of these last two objects is still uncertain. In Chapters 3 and 4 we studied possible scenarios to produce such long-period sources.

In particular in Chapter 3, by assuming a neutron-star nature for these sources, I showed that the standard dipolar losses invoked to explain the currently observed radio pulsar population are insufficient to reach such long spin periods unless extreme (constant) magnetic fields are taken into account. We therefore explored the possibility that such slow pulsars might be neutron stars spun down during an initial phase of fallback accretion from the supernova debris. Soon after the supernova, the matter that remains gravitationally bound to the central compact object can start to fallback and circularise to form a disk. This disk can interact with the central neutron star and exchange angular momentum depending primarily on the mass accretion rate and the magnetic-field strength of the central neutron star. I showed that newly born neutron stars with birth magnetic fields of $B_0 \sim 10^{14-15}$ G, experiencing a propeller accretion phase from a fallback disk with moderate initial accretion rates of 10^{22-27} g s^{-1}, can reach spin periods of thousands of seconds at relatively young ages ($\sim 10^{3-5}$ yr). On the other hand, neutron stars with lower magnetic fields are unaffected by the presence of a fallback disk and do not experience a significant spin-down. Hence, while predicting the existence of a population of neutron stars with very long periods, this model also explains the normal radio pulsar population. However deep X-ray observations have imposed upper limits on the X-ray luminosity of these two objects that challenge the magnetar interpretation. If they are neutron stars, they appear to be too cold to possess such strong magnetic fields according to our current understanding of magneto-thermal evolution in neutron stars.

In Chapter 4, we considered both the neutron-star and magnetic white-dwarf origin for these periodic sources and further studied the implications for the two scenarios. First, we showed that for both scenarios, these objects are difficult to reconcile with the standard model for radio emission due to particle acceleration in

magnetospheric gaps. Moreover, we performed population synthesis simulations of both neutron stars and magnetic white dwarfs to assess the possibility of having long-period sources with a sufficient amount of rotational energy to power the observed radio emission. Accounting for even the most extreme scenario of fallback accretion and constant magnetic fields, we showed that it is very unlikely to have a population of long-period neutron stars that still have the rotational energy budget to explain the observations of GLEAM-X J1627 and GPM J1839–10. On the other hand, since magnetic white dwarfs are born with longer spin periods and possess larger moments of inertia, a fraction of their population may possess the rotational energies and the magnetic fields necessary to power the observed emission.

More than 50 years after the first discovery of radio pulsations from neutron stars, the physical mechanism that produces coherent radio emission from rotating magnetosphere is still poorly understood. The discovery of PSR J0901-4046, GLEAM-X J1627 and GPM J1839–10 poses further challenges to our current understanding of neutron-star and white-dwarf emission and evolution. The long-period domain in the radio band has been poorly explored in the past and the fact that we serendipitously detected these sources indicates that they may be very common. Future dedicated Galactic plane imaging surveys focusing on long-period sources will likely discover many more objects of this kind. The theoretical ground work presented in this thesis will therefore provide useful insights into future discoveries and help to shed light on the nature of this sources.

Appendix A

Timing tests

A.1 Hardware and Software

Our test machine features an Intel(R) Xeon(R) Gold 6230R CPU at 2.10GHz with a single NVIDIA GTX 2080 Ti GPU, 16 GiB RAM, and SSD drives. The system is running CentOS Linux release 7.8.2003 (Core) with PyTorch 1.2.0, CUDA toolkit 10.0.130 and GPU driver 455.32.00.

A.2 Timing Tests

A.2.1 Timing for Single-parameter Predictions

We report here the timing benchmarks for the MLP and CNN during the single-parameter training experiments discussed in Section 5.4.1. We run our experiments on the test machine and individually record the forward pass time (the time needed to go through the samples in a batch and compute a prediction) and the backward pass time (the time to compute all the gradients and perform a single optimisation step) as a function of the batch size and resolution. Our benchmarks for the training data-sets from simulation run S1 using the four training configurations T1, T2, T3 and T4 are shown in Figures A.1 and A.2, respectively.

The timing benchmark shows that the MLP is slightly faster in performing an optimisation step. This is expected due to fewer trainable parameters when compared to the more complex CNN. We can also see that the forward and backward pass times per sample decrease with increasing batch size for both the MLP and CNN. For a larger batch size, several input samples are transferred from the CPU to the
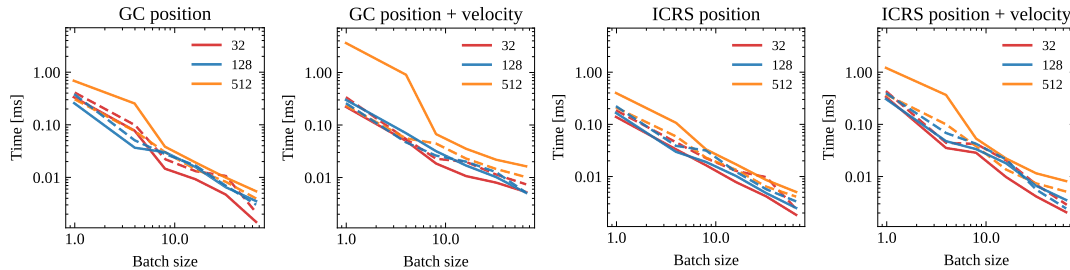
FIGURE A.1: MLP forward (*solid lines*) and backward (*dashed lines*) pass times per sample in ms for the training process on the single parameter σ_k of the Maxwell kick-velocity distribution, as a function of the batch size and the resolution (*red*, *blue*, and *orange* curves for 32, 128 and 512 respectively) using the four different input configurations T1 (GC position), T2 (GC position + velocity), T3 (ICRS position) and T4 (ICRS position + velocity).
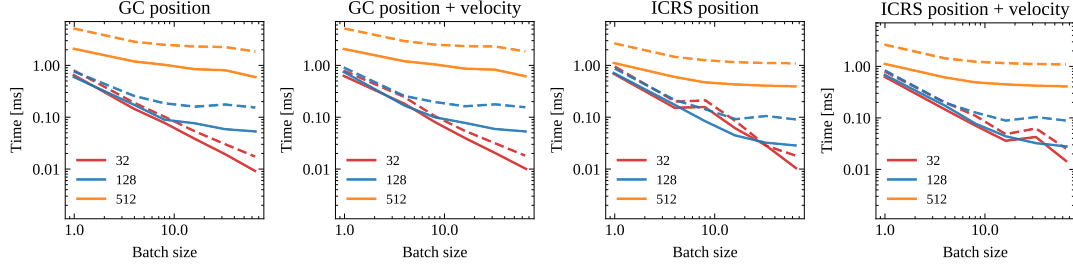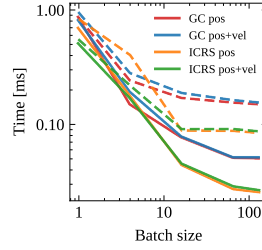
FIGURE A.2: CNN forward (*solid lines*) and backward (*dashed lines*) pass times per sample in ms for the training process on the single parameter σ_k of the Maxwell kick-velocity distribution, as a function of the batch size and the resolution (*red*, *blue*, and *orange* curves for 32, 128 and 512 respectively) using the four different input configurations T1 (GC position), T2 (GC position + velocity), T3 (ICRS position) and T4 (ICRS position + velocity).

FIGURE A.3: CNN forward pass (*solid line*) and backward pass time (*dashed line*) per sample for the two-parameter experiments as a function of the batch size for the four different input configurations T1 (GC position), T2 (GC position + velocity), T3 (ICRS position) and T4 (ICRS position + velocity).

GPU in one step, reducing the overall number of calls between the two. Thus, on average, the processing time for an individual sample reduces when the batch size is increased. Moreover, a higher resolution generally implies an increase in computational cost, albeit being more pronounced in the case of the CNN than the MLP. The number of input channels itself has very little effect on our timings. Finally, we note that ICRS maps are slightly faster to process (in particular for the higher resolutions), due to the fact that their size is smaller compared to the galactocentric maps (they have half the height in bins).

A.2.2 Timing for Two-parameter Predictions

Following the results for the single-parameter experiments, we restrict our two-parameter predictions to the CNN model only and fix the resolution of the input maps to 128. The results of our timing benchmarks using the training data-sets from simulation run S3 with the four configurations T1, T2, T3 and T4 are shown in Figure A.3. We again report the timings for the forward and backward passes per sample as a function of the batch size and the type of input channels provided. As for the single-parameter case, we conclude that using ICRS maps ensures the lowest forward and backward pass times.

Appendix B

Neutron stars with measured proper motion

In Table B.1, we report the properties of 417 neutron stars with measured proper motions in RA and DEC. Data for these neutron stars are primarily collected from the ATNF catalogue[1] (Manchester et al., 2005). In some cases, updated estimates are available and those values quoted and the corresponding references specified. Note that in those cases, where multiple proper-motion estimates are available, we choose the ones with the lowest absolute error. The columns report in order: (i) the object name based on J2000 coordinates; (ii) the right ascension (RA) in hour angle and (iii) declination (DEC) in degrees with the last digit uncertainty given in parentheses; the proper motion in (iv) RA and (v) DEC in milliarcseconds per year with corresponding uncertainties; (vi) the parallax measured in milliarcseconds with uncertainty where available; (vii) the position epoch in modified Julian days; (viii) the spin period in seconds; (ix) the spin-period derivative in seconds over seconds; (x) the dispersion measure in $[\mathrm{pc\,cm^{-3}}]$ with the last digit uncertainty given in parentheses; (xi) the heliocentric distance derived from the DM using the YMW16 free-electron density model (for some objects the DM exceeds the maximum Galactic DM allowed by the YMW16 model, which assigns a default value of 25 kpc; when available, we quote other distance estimates; * indicates a distance derived from other techniques especially for X-rays and gamma-ray sources, which have no measured DM); (xii) the classification of the object, i.e., radio pulsar (PSR), binary pulsar (binary PSR), gamma-/X-ray pulsar (Gamma-/X-ray PSR), magnetar (MAG), X-ray dim isolated neutron star (XDINS); if the object is associated with a globular cluster (GC) or the Small Magellanic Cloud (SMC) this is reported in brackets; and (xiii) the reference for the proper motion measurements, indicated only if different from the ATNF catalogue, i.e., [1] Motch et al. (2009), [2] Eisenbeiss et al. (2010), [3] Walter et al. (2010), [4] Stovall et al. (2014), [5] Jennings et al. (2018), [6] Perera et al. (2019), [7] Dang et al. (2020), [8] Danilenko et al. (2020).

---

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

TABLE B.1: Up-to-date list of 417 neutron stars with measured proper motions in RA and DEC.

JName	RA [hms]	DEC [° ′ ″]	μ_{RA} [mas yr⁻¹]	μ_{DEC} [mas yr⁻¹]	parallax [mas]	Pos. Epoch [MJD]	P [s]	$\dot{P}$ [s s⁻¹]	DM [pc cm⁻³]	$d_\odot$ [kpc]	Class(Assoc.)	Ref.
J0014+4746	00:14:17.75(4)	+47:46:33.4(3)	19.3±1.8	−19.7±1.5	—	49664	1.24070	5.6446×10^{-16}	30.405(13)	1.776	PSR	
J0023−0923	00:23:16.87910(2)	+09:23:23.8689(8)	−12.4±0.5	−6.1±0.1	0.4±0.3	56567	0.00305	1.14234×10^{-20}	14.32810(4)	1.248	PSR	[6]
J0024−7204ab	00:24:08.1615(5)	−72:04:47.602(2)	4.2±0.6	−2.9±0.5	—	51600	0.00370	9.820×10^{-21}	24.37(2)	2.54	PSR(GC)	
J0024−7204C	00:23:50.3546(1)	−72:04:31.5048(4)	5.2±0.1	−3.1±0.1	—	51600	0.00576	-4.98503×10^{-20}	24.5955(8)	2.594	PSR(GC)	
J0024−7204D	00:24:13.88092(6)	−72:04:43.8524(2)	4.24±0.07	−2.24±0.05	—	51600	0.00536	-3.4220×10^{-21}	24.7432(9)	2.631	PSR(GC)	
J0024−7204E	00:24:11.10528(5)	−72:05:20.1492(2)	6.15±0.03	−2.35±0.06	—	51600	0.00354	9.85103×10^{-20}	24.236(2)	2.509	PSR(GC)	
J0024−7204F	00:24:03.85547(10)	−72:04:42.8183(2)	4.52±0.08	−2.5±0.05	—	51600	0.00262	6.45029×10^{-20}	24.382(5)	2.543	PSR(GC)	
J0024−7204G	00:24:07.9603(1)	−72:04:39.7030(5)	4.5±0.1	−2.9±0.1	—	51600	0.00404	-4.21584×10^{-20}	24.436(4)	2.556	PSR(GC)	
J0024−7204H	00:24:06.7032(2)	−72:04:06.8067(6)	5.1±0.2	−2.8±0.2	—	51600	0.00321	-1.8294×10^{-21}	24.369(8)	2.541	PSR(GC)	
J0024−7204I	00:24:07.9347(2)	−72:04:39.6815(7)	5±0.2	−2.1±0.2	—	51600	0.00348	-4.5874×10^{-20}	24.43(1)	2.555	PSR(GC)	
J0024−7204J	00:23:59.4077(1)	−72:03:58.7908(5)	5.27±0.06	−3.59±0.09	—	51600	0.00210	-9.7917×10^{-21}	24.5932(1)	2.594	PSR(GC)	
J0024−7204L	00:24:03.7721(3)	−72:04:56.923(2)	4.4±0.2	−2.4±0.2	—	51600	0.00435	-1.22045×10^{-19}	24.40(1)	2.547	PSR(GC)	
J0024−7204M	00:23:54.4899(3)	−72:05:30.756(2)	5±0.3	−2±0.4	—	51600	0.00368	-3.8419×10^{-20}	24.43(2)	2.553	PSR(GC)	
J0024−7204N	00:24:09.1880(2)	−72:04:28.8907(7)	6.3±0.2	−2.8±0.2	—	51600	0.00305	-2.18570×10^{-20}	24.574(9)	2.589	PSR(GC)	
J0024−7204O	00:24:04.65254(6)	−72:04:53.7670(2)	5.01±0.05	−2.58±0.08	—	51600	0.00264	3.03493×10^{-20}	24.356(2)	2.537	PSR(GC)	
J0024−7204Q	00:24:16.4909(1)	−72:04:25.1644(6)	5.2±0.1	−2.6±0.1	—	51600	0.00403	3.4008×10^{-20}	24.265(4)	2.517	PSR(GC)	
J0024−7204R	00:24:07.5851(2)	−72:04:50.3954(5)	4.8±0.1	−3.3±0.2	—	51600	0.00348	1.48352×10^{-19}	24.361(7)	2.538	PSR(GC)	
J0024−7204S	00:24:03.9794(1)	−72:04:42.3530(4)	4.5±0.1	−2.5±0.1	—	51600	0.00283	$-1.205413 \times 10^{-19}$	24.376(4)	2.542	PSR(GC)	
J0024−7204T	00:24:08.5491(5)	−72:04:38.932(3)	5.1±0.6	−2.6±0.7	—	51600	0.00759	2.93805×10^{-19}	24.41(2)	2.55	PSR(GC)	
J0024−7204U	00:24:09.8366(1)	−72:03:59.6882(4)	4.6±0.2	−3.8±0.1	—	51600	0.00434	9.52279×10^{-20}	24.337(4)	2.534	PSR(GC)	
J0024−7204W	00:24:06.058(1)	−72:04:49.088(2)	6.1±0.5	−2.6±0.3	—	50000	0.00235	-8.65526×10^{-20}	24.367(3)	2.539	PSR(GC)	
J0024−7204X	00:24:22.38565(9)	−72:01:17.4414(7)	5.8±0.1	−3.3±0.2	—	54000	0.00477	1.83610×10^{-20}	24.539(5)	2.587	PSR(GC)	
J0024−7204Y	00:24:01.4026(1)	−72:04:41.8363(4)	4.4±0.1	−3.4±0.1	—	51600	0.00220	-3.51721×10^{-20}	24.468(4)	2.563	PSR(GC)	
J0024−7204Z	00:24:06.041(2)	−72:05:01.480(6)	4±2	1±2	—	51600	0.00455	-4.54×10^{-21}	24.45(4)	2.559	PSR(GC)	
J0030+0451	00:30:27.42838(3)	+04:51:39.707(1)	−6.4±0.1	0.9±0.3	2.94±0.09	55390	0.00487	1.01743×10^{-20}	4.33991(7)	0.345	PSR	[6]
J0034−0534	00:34:21.834228(8)	−05:34:36.722(3)	7.9±0.3	−9.2±0.6	—	55000	0.00188	4.966×10^{-21}	13.76517(4)	1.348	PSR	[6]
J0034−0721	00:34:08.8703(1)	−07:21:53.409(2)	10.37±0.08	−11.13±0.16	0.93±0.08	52275	0.94295	4.08210×10^{-16}	10.922(6)	0.996	PSR	
J0040+5716	00:40:32.3899(1)	+57:16:24.833(1)	12.39±0.05	−5.4±0.1	0.11±0.04	49667	1.11823	2.878973×10^{-15}	92.5146(25)	2.42	PSR	
J0045−7319	00:45:35.255(10)	−73:19:03.32(1)	0.31±0.11	−0.931±0.099	0.04±0.06	57205.875	0.92628	4.4632×10^{-15}	105.4(7)	58.725	Binary PSR(SMC)	[5]
J0055+5117	00:55:45.3981(1)	+51:17:24.601(1)	10.47±0.07	−17.4±0.1	0.35±0.06	49664	2.11517	9.53764×10^{-15}	44.0127(24)	1.935	PSR	
J0101−6422	01:01:11.1163(2)	−64:22:30.171(2)	10±1	−12±2	—	55520	0.00257	5.16×10^{-21}	11.926(1)	1.002	PSR	
J0102+6537	01:02:32.9914(1)	+65:37:13.416(1)	9.24±0.07	1.77±0.15	0.4±0.04	49675	1.67916	5.95492×10^{-15}	65.853(16)	1.978	PSR	
J0108−1431	01:08:08.347016(88)	−14:31:50.1871(11)	75.1±2.3	−152.5±1.7	4.2±1.4	54100	0.80756	7.704×10^{-17}	2.38(19)	0.232	PSR	
J0108+6608	01:08:22.5049(3)	+66:08:34.499(1)	−32.75±0.03	35.15±0.04	0.47±0.03	50011	1.28366	1.3051×10^{-14}	30.5482(14)	1.461	PSR	
J0134−2937	01:34:18.6824(3)	−29:37:17.042(7)	17±1	−9±2	—	54428	0.13696	7.83699×10^{-17}	21.806(6)	25	PSR	
J0139+5814	01:39:19.7401(12)	+58:14:31.819(17)	−19.11±0.07	−16.6±0.07	0.37±0.04	52275	0.27245	1.071167×10^{-14}	73.81141(76)	2.006	PSR	

JName	RA [hms]	DEC [° ′ ″]	μ_{RA} [mas yr⁻¹]	μ_{DEC} [mas yr⁻¹]	parallax [mas]	Pos. Epoch [MJD]	P [s]	$\dot{P}$ [s s⁻¹]	DM [pc cm⁻³]	$d_\odot$ [kpc]	Class(Assoc.)	Ref.
J0146+6145	01:46:22.42(10)	+61:45:02.8(7)	−4.1±1	1.9±1	—	49031	8.68899	1.99×10^{-12}	—	3.6*	MAG	
J0147+5922	01:47:44.6434(1)	+59:22:03.284(1)	−6.38±0.092	3.826±0.075	0.495±0.06	53160	0.19632	2.56763×10^{-16}	40.111(3)	1.578	PSR	
J0151−0635	01:51:22.7179(2)	−06:35:02.987(2)	10.697±0.119	−5.373±0.069	0.217±0.08	49347	1.46466	4.4259×10^{-16}	25.66(3)	25	PSR	
J0152−1637	01:52:10.8539(1)	−16:37:53.641(2)	0.804±0.217	−31.372±0.368	0.44±0.25	48227	0.83274	1.299202×10^{-15}	11.92577(4)	0.92	PSR	
J0154+1833	01:54:36.88273(3)	+18:33:50.758(9)	10.3±0.9	−8.9±1.9	—	56900	0.00236	2.919×10^{-21}	19.7978(1)	1.622	PSR	
J0157+6212	01:57:49.9434(1)	+62:12:26.648(1)	1.57±0.06	44.8±0.04	0.56±0.03	49709	2.35174	1.889254×10^{-13}	30.21(3)	1.385	PSR	
J0205+6449	02:05:37.92247(15)	+64:49:41.3343(1)	−1.4±0.16	0.54±0.16	—	55487	0.06572	$1.93754256 \times 10^{-13}$	140.7(3)	2.784	PSR	[6]
J0206−4028	02:06:01.293(5)	−40:28:03.62(7)	9±4	26±5	—	57600	0.63055	1.1966×10^{-15}	12.9(12)	1.26	PSR	
J0218+4232	02:18:06.358569(4)	+42:32:17.37515(8)	5.31±0.08	−3.1±0.1	0.16±0.09	55849.1	0.00232	7.73955×10^{-20}	61.252(5)	2.928	PSR	[6]
J0248+6021	02:48:18.617(1)	+60:21:34.72(1)	48±10	48±4	—	54000	0.21709	5.509×10^{-14}	370(1)	2*	PSR	
J0255−5304	02:55:56.2925(9)	−53:04:21.275(9)	25±4	7±3	—	57600	0.44771	3.129×10^{-17}	15.9(5)	1.514	PSR	
J0304+1932	03:04:33.115(16)	+19:32:51.4(8)	6±7	−37±4	—	49289	1.38758	1.295235×10^{-15}	15.65677(35)	0.739	PSR	
J0323+3944	03:23:26.6619(1)	+39:44:52.403(1)	26.5±0.05	−30.77±0.02	1.05±0.04	49290	3.03207	6.3562×10^{-16}	26.18975(93)	1.197	PSR	
J0332+5434	03:32:59.4096(1)	+54:34:43.329(1)	16.97±0.03	−10.37±0.05	0.59±0.02	46473	0.71452	2.048265×10^{-15}	26.7641(1)	1.183	PSR	
J0335+4555	03:35:16.6416(1)	+45:55:53.452(1)	−3.66±0.05	−0.08±0.12	0.41±0.02	49912	0.26920	7.3505×10^{-18}	47.14571(29)	1.535	PSR	
J0337+1715	03:37:43.82589(13)	+17:15:14.828(2)	4.81±0.49	−4.42±0.42	0.73±0.24	55920	0.00273	1.7666×10^{-20}	21.3162(3)	0.816	Binary PSR	
J0340+4130	03:40:23.28822(61)	41:30:45.2900(3)	0±0.4	−4.4±0.8	0.4±0.4	56675	0.00330	7.0470×10^{-21}	49.5865(16)	1.603	PSR	[6]
J0348+0432	03:48:43.639000(4)	+04:32:11.4580(2)	4.04±0.16	3.5±0.6	—	56000	0.03912	2.4075×10^{-19}	40.46313(11)	2.256	Binary PSR	
J0357+3205	03:57:52.5(2)	+32:05:25(6)	117±20	115±20	—	54946	0.44410	1.3037×10^{-14}	—	0.83*	Gamma-/X-ray PSR	
J0357+5236	03:57:44.8403(2)	+52:36:57.493(1)	13.908±0.088	−10.633±0.078	0.305±0.05	49912	0.19703	4.76594×10^{-16}	103.706(4)	2.025	PSR	
J0358+5413	03:58:53.7238(2)	+54:13:13.784(3)	9.3±0.5	8.3±0.8	0.91±0.16	54096	0.15638	4.39496×10^{-15}	57.1420(3)	1.594	PSR	
J0406+6138	04:06:30.0806(2)	+61:38:41.408(1)	12.4±0.1	22.74±0.08	0.22±30.05	53784	0.59457	5.57316×10^{-15}	65.4053(34)	1.785	PSR	
J0437−4715	04:37:15.8961737(6)	−47:15:09.110714(7)	121.443±0.001	−71.474±0.002	6.42±0.07	54500	0.00576	5.729215×10^{-20}	2.64476(7)	0.156	Binary PSR	[6]
J0452−1759	04:52:34.1057(1)	−17:59:23.371(2)	11.5±0.9	6±2	0.6±1.4	52275	0.54894	5.753107×10^{-15}	39.903(3)	2.711	PSR	
J0453+1559	04:53:45.41368(5)	+15:59:21.3063(59)	−5.5±0.5	−6±4.2	—	56400	0.04578	1.8612×10^{-19}	30.3053(3)	0.522	PSR	
J0454+5543	04:54:07.7506(1)	+55:43:41.437(2)	53.34±0.06	−17.56±0.14	0.84±0.05	52275	0.34073	2.37227×10^{-15}	14.59002(15)	0.637	PSR	
J0502+4654	05:02:04.561(4)	+46:54:06.09(9)	−8±3	8±5	—	48717	0.63857	5.582339×10^{-15}	41.834(20)	1.319	PSR	
J0509+0856	05:09:22.23485(1)	+08:56:25.02488(8)	5.4±0.2	−4.3±0.5	—	57384	0.00406	4.412×10^{-21}	38.318(4)	0.817	PSR	
J0514−4002A	05:14:06.69271(2)	−40:02:48.8930(19)	5.19±0.22	−0.56±0.25	—	53623.16	0.00499	7.0×10^{-22}	52.14016(37)	13.2*	PSR(GC)	[7]
J0520−2553	05:20:36.185(3)	−25:53:12.28(1)	20±12	19±21	—	51216	0.24164	3.0092×10^{-17}	33.77(3)	2.379	PSR	
J0525+1115	05:25:56.4483(11)	+11:15:19.08(7)	31±2	−1±8	—	48262	0.35444	7.36052×10^{-17}	79.418(13)	1.841	PSR	
J0528+2200	05:28:52.264(9)	+22:00:04(2)	−20±19	7±9	—	53100	3.74554	4.00530×10^{-14}	50.8695(13)	1.215	PSR	
J0534+2200	05:34:31.973(5)	+22:00:52.06(6)	−14.7±0.8	2±0.8	—	40706	0.03339	4.20972×10^{-13}	56.77118(24)	1.311	PSR	
J0536−7543	05:36:30.82(1)	−75:43:54.8(1)	3±40	65±8	—	57600	1.24586	5.777×10^{-16}	18.58(2)	0.143	PSR	
J0538+2817	05:38:25.0572(1)	+28:17:09.161(2)	−23.57±0.1	52.87±0.1	0.72±0.12	52275	0.14316	3.694515×10^{-15}	39.570(1)	0.946	PSR	
J0543+2329	05:43:11.00(17)	+23:16:40.10(47)	19±7	12±8	—	55580	0.24598	1.541956×10^{-14}	77.7026(10)	1.565	PSR	
J0601−0527	06:01:58.9752(2)	−05:27:50.871(4)	−7.36±0.07	−15.24±0.09	0.48±0.04	49379	0.39597	1.302098×10^{-15}	80.538(5)	2.328	PSR	
J0610−2100	06:10:13.595462(17)	−21:00:27.9313(4)	9.04±0.08	16.73±0.08	—	55000	0.00386	$0.00001235 \times 10^{-15}$	60.6662(17)	3.259	PSR	[6]

JName	RA [hms]	DEC [° ′ ″]	μ_{RA} [mas yr^{-1}]	μ_{DEC} [mas yr^{-1}]	parallax [mas]	Pos. Epoch [MJD]	P [s]	$\dot{P}$ [s s^{-1}]	DM [pc cm^{-3}]	$d_{\odot}$ [kpc]	Class(Assoc.)	Ref.
J0613-0200	06:13:43.975698(10)	-02:00:47.22547(3)	1.828 ± 0.005	-10.35 ± 0.01	0.9 ± 0.04	55413	0.00306	9.59057 × 10^{-21}	38.77919(6)	1.024	PSR	[6]
J0614+2229	06:14:17.0058(1)	+22:29:56.848(1)	-0.24 ± 0.04	-1.25 ± 0.04	0.28 ± 0.03	49674	0.33496	5.94494 × 10^{-14}	96.91(4)	1.744	PSR	
J0614-3329	06:14:10.3478(3)	-33:29:54.118(4)	0.58 ± 0.09	-1.92 ± 0.12	1.1 ± 1.1	55100	0.00315	1.75 × 10^{-20}	37.049(1)	2.69	PSR	[6]
J0621+1002	06:21:22.11436(3)	+10:02:38.7352(15)	3.2 ± 0.1	-0.6 ± 0.5	—	55000	0.02885	4.730 × 10^{-20}	36.47(3)	0.42	PSR	
J0629+2415	06:29:05.7273(1)	+24:15:41.546(1)	3.56 ± 0.12	-4.68 ± 0.08	0.33 ± 0.05	49438	0.47662	1.99573 × 10^{-15}	84.1762(50)	1.667	PSR	
J0630-2834	06:30:49.404393(43)	-28:34:42.77881(37)	-46.3 ± 0.99	21.26 ± 0.52	3.01 ± 0.41	54100	1.24442	7.1229 × 10^{-15}	34.425(1)	2.069	PSR	
J0633-0632	06:33:44.21(2)	06:32:34.9(16)	-36 ± _	39 ± _	—	_	0.29740	7.9592 × 10^{-14}	_	1.35*	Gamma-/X-ray PSR	[8]
J0633+1746	06:33:54.1530(28)	+17:46:12.909(4)	138 ± 4	97 ± 4	6.4 ± 1.8	49793	0.23710	1.097087 × 10^{-14}	2.89(2)	0.138	PSR	[6]
J0636+5129	06:36:04.8472(39)	+51:28:59.967(82)	4.3 ± 0.9	2 ± 1.2	4.9 ± 0.6	57002	0.00287	3.442 × 10^{-21}	11.10816(11)	0.21	PSR	[4]
J0645+5158	06:45:59.0821(2)	+51:58:14.919(41)	1.64 ± 0.08	-7.2 ± 0.1	0.39 ± 0.08	56534	0.00885	4.9219 × 10^{-21}	18.247536(9)	0.672	PSR	[6]
J0653+8051	06:53:15.09(3)	+80:52:00.22(11)	19 ± 3	-1 ± 3	—	48712	1.21444	3.798787 × 10^{-15}	33.31882(71)	2.348	PSR	
J0659+1414	06:59:48.196(3)	+14:14:19.4(3)	44.07 ± 0.63	-2.4 ± 0.29	3.47 ± 0.36	57600	0.38493	5.494123 × 10^{-14}	13.94(9)	0.159	PSR	
J0709+0458	07:09:08.36539(2)	+04:58:51.4941(6)	-0.7 ± 0.2	-1.3 ± 0.4	—	56983	0.03443	3.8040 × 10^{-19}	44.2664(5)	1.203	PSR	
J0711-6830	07:11:54.189114(13)	-68:30:47.41446(8)	-15.57 ± 0.02	14.24 ± 0.02	—	54500	0.00549	1.49067 × 10^{-20}	18.4066(6)	0.106	PSR	[6]
J0720-3125	07:20:24.9620(9)	-31:25:50.083(11)	-93.2 ± 5.4	46.8 ± 5.1	2.8 ± 0.9	52645.1	8.39112	6.983 × 10^{-14}	_	0.4*	XDINS	[2]
J0729-1836	07:29:32.3369(1)	-18:36:42.244(2)	-13.06 ± 0.11	13.27 ± 0.44	0.5 ± 0.09	49720	0.51016	1.895713 × 10^{-14}	61.293(10)	2.403	PSR	
J0737-3039A	07:37:51.248419(26)	-30:39:40.71431(1)	-3.82 ± 0.62	2.13 ± 0.23	0.87 ± 0.14	54100	0.02270	1.75993 × 10^{-18}	48.920(5)	1.105	PSR	
J0737-3039B	07:37:51.248419(26)	-30:39:40.71431(1)	-3.82 ± 0.62	2.13 ± 0.23	0.87 ± 0.14	54100	2.77346	8.92 × 10^{-16}	48.920(5)	1.105	PSR	
J0738-4042	07:38:32.262(2)	-40:42:39.42(2)	-14 ± 1.2	13 ± 2	—	57600	0.37492	1.37504 × 10^{-15}	160.896(3)	1.57	PSR	
J0742-2822	07:42:49.058(2)	-28:22:43.76(4)	-29 ± 2	4 ± 2	—	49326	0.16676	1.682112 × 10^{-14}	73.728(1)	3.114	PSR	
J0745-5353	07:45:04.48(4)	-53:53:09.56(3)	-60 ± 10	50 ± 10	—	55129	0.21484	2.18686 × 10^{-15}	121.38(2)	0.565	PSR	
J0751+1807	07:51:09.155331(13)	+18:07:38.4864(1)	-2.72 ± 0.06	-13.4 ± 0.3	0.8 ± 0.2	55000	0.00348	0.77874 × 10^{-20}	30.246(6)	0.428	PSR	[6]
J0754+3231	07:54:40.688(5)	+32:31:56.2(2)	-4 ± 5	7 ± 3	—	48725	1.44235	1.079865 × 10^{-15}	39.9863(14)	1.464	PSR	
J0758-1528	07:58:29.0708(7)	-15:28:08.738(19)	-10 ± 4	12 ± 3	—	49896	0.68227	1.618890 × 10^{-15}	63.327(3)	2.71	PSR	
J0809-4753	08:09:43.835(6)	-47:53:54.70(6)	-19 ± 5	15 ± 4	—	57600	0.54720	3.0763 × 10^{-15}	228.3(12)	6.488	PSR	
J0814+7429	08:14:59.50(2)	+74:29:05.70(11)	24.02 ± 0.09	-44 ± 0.4	2.31 ± 0.04	49162	1.29224	1.68114 × 10^{-16}	5.75066(48)	0.368	PSR	
J0820-1350	08:20:26.3817(1)	-13:50:55.859(2)	21.64 ± 0.09	-39.44 ± 0.05	0.51 ± 0.04	52275	1.23813	2.105178 × 10^{-15}	40.938(3)	2.321	PSR	
J0821-4300	08:21:57.3653(31)	-43:00:17.074(43)	-54.1 ± 8.3	-28.1 ± 10	—	53454	0.11280	9.3 × 10^{-18}	_	2.2*	Gamma-/X-ray PSR	
J0823+0159	08:23:09.7651(1)	+01:59:12.469(1)	-3.797 ± 0.244	0.171 ± 0.256	0.376 ± 0.09	49281	0.86487	1.04550 × 10^{-16}	23.727(6)	0.808	PSR	
J0824+0028	08:24:24.84022(5)	+00:28:00.567(2)	-4.3 ± 0.4	-9.2 ± 1.3	—	56600	0.00986	1.46890 × 10^{-19}	34.5485(6)	1.689	PSR	
J0826+2637	08:26:51.5068(1)	+26:37:21.297(1)	63 ± 0.01	-96.75 ± 0.07	2.01 ± 0.01	46450	0.53066	1.70932 × 10^{-15}	19.47633(18)	0.314	PSR	
J0835-4510	08:35:20.61149(2)	-45:10:34.8751(3)	-49.68 ± 0.06	29.9 ± 0.1	3.5 ± 0.2	51544	0.08933	1.25008 × 10^{-13}	67.97(2)	0.328	PSR	
J0837-0610	08:37:05.642(3)	+06:10:14.56(14)	2 ± 5	51 ± 3	—	48721	1.27377	6.799219 × 10^{-15}	12.8640(4)	0.19	PSR	
J0837+4135	08:37:21.1943(6)	-41:35:14.563(6)	11 ± 2	-13 ± 1	—	57600	0.75163	3.53848 × 10^{-15}	147.29(7)	1.426	PSR	
J0846-3533	08:46:06.040(8)	-35:33:41.2(1)	93 ± 72	-15 ± 65	—	57600	1.11610	1.6044 × 10^{-15}	94.16(11)	0.535	PSR	
J0900-3144	09:00:43.953088(8)	-31:44:30.89520(13)	-1.01 ± 0.05	2.01 ± 0.06	0.7 ± 0.6	55000	0.01111	4.8880 × 10^{-20}	75.707(8)	0.38	PSR	[6]
J0908-1739	09:08:38.1822(11)	-17:39:37.67(3)	27 ± 11	-40 ± 11	—	48737	0.40163	6.695000 × 10^{-16}	15.879(2)	0.805	PSR	
J0922+0638	09:22:14.022(4)	+06:38:23.30(17)	18.8 ± 0.9	86.4 ± 0.7	0.83 ± 0.13	48227	0.43063	1.37294 × 10^{-14}	27.2986(5)	1.912	PSR	[6]

JName	RA [hms]	DEC [° ′ ″]	μ_{RA} [mas yr^{-1}]	μ_{DEC} [mas yr^{-1}]	parallax [mas]	Pos. Epoch [MJD]	P [s]	$\dot{P}$ [s s^{-1}]	DM [pc cm^{-3}]	$d_\odot$ [kpc]	Class(Assoc.)	Ref.
J0942-5657	09:42:54.424(9)	-56:57:43.21(6)	-12±7	12±3	–	57600	0.80816	3.96051×10^{-14}	159.74(12)	0.406	PSR	
J0943-1631	09:43:30.10(15)	+16:31:37(6)	23±16	9±11	–	48865	1.08742	9.109×10^{-17}	20.3402(18)	1.487	PSR	
J0944-1354	09:44:28.9559(11)	-13:54:41.63(3)	-1±32	-22±14	–	49337	0.57026	4.5256×10^{-17}	12.497(3)	0.686	PSR	
J0946-0951	09:46:07.6(8)	+09:51:55(34)	-38±19	-21±12	–	48483	1.09771	3.49339×10^{-15}	15.31845(90)	0.892	PSR	
J0953-0755	09:53:09.3097(19)	+07:55:35.75(8)	-2.09±0.08	29.46±0.07	3.82±0.07	46375	0.25307	2.29758×10^{-16}	2.96927(8)	0.187	PSR	
J1012-5307	10:12:33.4382(63)	+53:07:02.283(77)	2.61±0.01	-25.49±0.01	0.9±0.2	55291	0.00526	1.712794×10^{-20}	9.02314(7)	0.805	Binary PSR	[6]
J1017-7156	10:17:51.32828(2)	-71:56:41.64586(11)	-7.31±0.06	6.76±0.05	3.9±1.2	55329.1	0.00234	2.2194×10^{-21}	94.22407(3)	1.807	PSR	
J1022+1001	10:22:57.9957(1)	+10:01:52.765(2)	-14.912±0.041	5.611±0.034	1.387±0.03	55000	0.01645	4.33375×10^{-20}	10.2521(1)	0.834	Binary PSR	[6]
J1023+0038	10:23:47.68720(3)	+00:38:40.8455(5)	4.76±0.03	-17.34±0.04	0.731±0.02	54906	0.00169	6.930×10^{-21}	14.325(10)	1.112	Binary PSR	[6]
J1024-0719	10:24:38.6673(34)	-07:19:19.595(11)	-35.27±0.02	-48.22±0.04	0.8±0.1	56239	0.00516	1.85558×10^{-20}	6.4778(6)	0.381	Binary PSR	[6]
J1035-6720	10:35:27.478(1)	-67:20:12.692(6)	-12±3	1±3	–	55716	0.00287	4.6463×10^{-20}	84.16(22)	1.46	PSR	
J1041-1942	10:41:36.196(5)	-19:42:13.61(11)	-1±3	14±5	–	48738	1.38637	9.44850×10^{-16}	33.777(10)	2.533	PSR	
J1045-4509	10:45:50.18696(3)	-45:09:54.1223(4)	-6.08±0.06	5.15±0.6	2.2±1.1	54500	0.00747	1.7662×10^{-20}	58.1438	0.336	PSR	[6]
J1047-3032	10:47:00.815(6)	-30:32:18.00(1)	55±54	34±69	–	51019	0.33033	6.1006×10^{-17}	52.54(7)	2.144	PSR	[7]
J1048+2339	10:48:43.4183(2)	+23:39:53.411(7)	-16.3±1	-11.7±1.2	0.96±0.8	57205.875	0.00467	3.003×10^{-20}	16.6544(1)	2.002	Binary PSR	[5]
J1057-5226	10:57:59.045(2)	-52:26:56.33(2)	42±5	-3±5	–	57600	0.19711	5.83870×10^{-15}	29.69(1)	0.093	PSR	
J1115+5030	11:15:38.400(5)	+50:30:12.29(5)	22±3	-51±3	–	49334	1.65644	2.492785×10^{-15}	9.18634(26)	0.923	PSR	
J1116-4122	11:16:43.086(1)	-41:22:43.96(1)	-1±5	7±20	–	51360	0.94316	7.9536×10^{-15}	40.53(3)	0.28	PSR	
J1121-5444	11:21:19.09(1)	-54:44:05.1(1)	-17±6	18±6	–	57600	0.53579	2.7840×10^{-15}	204.7(6)	2.324	PSR	
J1125-5825	11:25:44.36564(5)	-58:25:16.8798(4)	-10±0.3	2.4±0.3	–	55126.3	0.00310	6.08990×10^{-20}	124.7946(8)	1.744	PSR	
J1136-1551	11:36:03.1198(1)	+15:51:14.183(1)	-73.785±0.02	366.559±0.06	2.687±0.02	46407	1.18791	3.733837×10^{-15}	4.84066(34)	0.414	PSR	[6]
J1136-5525	11:36:02.16(2)	-55:25:06.9(2)	10±17	43±9	–	57600	0.36471	8.2266×10^{-15}	85.5(7)	1.522	PSR	
J1141-6545	11:41:07.0140(2)	-65:45:19.1131(15)	2.9±1	1.8±0.6	–	54637	0.39390	4.30737×10^{-15}	116.080(1)	1.676	PSR	
J1142+0119	11:42:50.145(3)	+01:19:35.69(9)	-41±29	-97±66	–	56315	0.00508	14.99×10^{-21}	19.197(1)	2.169	PSR	
J1202-5820	12:02:28.325(5)	-58:20:33.64(5)	-14±4	1±5	–	57600	0.45280	2.13023×10^{-15}	145.41(19)	1.887	PSR	
J1224-6407	12:24:22.254(6)	-64:07:53.87(4)	20±4	2±2	–	55191	0.21648	4.95328×10^{-15}	97.686(4)	1.535	PSR	
J1227-4853	12:27:58.724(1)	-48:53:42.741(3)	-18.73±0.2	7.39±0.11	0.62±0.16	56707.98	0.00169	1.11×10^{-20}	43.4235(7)	1.244	Binary PSR	[5]
J1231-1411	12:31:11.3132(7)	-14:11:43.63(2)	-62.03±0.26	6.2±0.5	–	55100	0.00368	2.280×10^{-20}	8.090(1)	0.42	PSR	
J1239+2453	12:39:40.4614(12)	+24:53:49.29(2)	-104.5±1.1	49.4±1.4	1.16±0.08	46531	1.38245	9.60049×10^{-16}	9.25159(53)	0.827	PSR	
J1253-5820	12:53:28.342(2)	-58:20:40.89(2)	24±2	13±2	–	49020	0.25550	2.09963×10^{-15}	100.584(4)	1.641	PSR	
J1257-1027	12:57:04.7625(2)	-10:27:05.551(2)	-7.164±0.122	12.079±0.114	0.141±0.07	49667	0.61731	3.62700×10^{-16}	29.634(9)	25	PSR	
J1300+1240	13:00:03.5767(1)	+12:40:56.4721(3)	45.5±0.05	-84.7±0.07	1.41±0.08	49750	0.00622	1.143341×10^{-19}	10.16550(3)	0.877	PSR	
J1302-6350	13:02:47.6542(1)	-63:50:08.6259(1)	-6.986±0.043	-0.416±0.044	0.418±0.03	57205.875	0.04776	2.2787×10^{-15}	146.73(1)	2.214	Binary PSR	[5]
J1308+2127	13:08:48.27(1)	+21:27:06.78(5)	-207±20	84±20	–	53415.68	10.3125	1.120×10^{-13}	–	25	XDINS	[1]
J1311-1228	13:11:52.656(9)	-12:28:01.1(3)	-2±8	-20±26	–	54371	0.44752	1.50859×10^{-16}	36.214(8)	25	PSR	
J1312+0051	13:12:46.64232(9)	+00:51:00.104(3)	-22.35±0.66	-11.2±1.5	–	55600	0.00423	17.527×10^{-21}	15.345(1)	1.471	PSR	
J1320-5359	13:20:53.921(3)	-53:59:05.00(4)	8±4	49±4	–	57600	0.27974	9.24577×10^{-15}	97.1(1)	2.204	PSR	[7]
J1321+8323	13:21:45.6315(7)	+83:23:39.432(1)	-52.674±0.087	32.373±0.126	0.968±0.08	48889	0.67004	5.6630×10^{-16}	13.31624(76)	0.977	PSR	

JName	RA [hms]	DEC [° ′ ″]	μ_{RA} [mas yr⁻¹]	μ_{DEC} [mas yr⁻¹]	parallax [mas]	Pos. Epoch [MJD]	P [s]	$\dot{P}$ [s s⁻¹]	DM [pc cm⁻³]	$d_\odot$ [kpc]	Class(Assoc.)	Ref.
J1327-0755	13:27:57.5880(12)	-07:55:29.80(4)	27±18	95±48	–	55131	0.00268	1.773×10^{-21}	27.91215(6)	25	PSR	
J1328-4357	13:28:06.432(1)	-43:57:44.12(1)	3±7	54±23	–	51360	0.53270	3.014×10^{-15}	42.0(10)	1.424	PSR	
J1405-4656	14:05:21.4255(8)	-46:56:02.31(1)	-44±6	20±10	–	55692.7	0.00760	2.79×10^{-20}	13.884(3)	0.669	PSR	
J1411+2551	14:11:18.866(3)	+25:51:08.39(7)	-3±12	-4±9	0.221±0.07	57617	0.06245	9.56×10^{-20}	12.3737(3)	1.131	Binary PSR	
J1417-4402	14:17:30.60(9)	-44:02:57.4(9)	-4.7±0.1	-5.1±0.87	–	57022	0.00266	–	55.00(3)	2.163	Binary PSR	
J1421-4409	14:21:20.9646(3)	-44:09:04.541(4)	-10±8	3±2	–	57600	0.00639	1.27×10^{-20}	54.635(4)	2.085	PSR	
J1430-6623	14:30:40.735(2)	-66:23:05.55(1)	-31±5	-21±3	–	57600	0.78544	2.77752×10^{-15}	65.3(1)	1.33	PSR	
J1431-4715	14:31:44.6177(2)	-47:15:27.574(4)	-12.01±0.33	14.51±0.26	0.64±0.16	55756.1	0.00201	1.411×10^{-20}	59.35(1)	1.822	Binary PSR	
J1446-4701	14:46:35.71391(2)	-47:01:26.7675(4)	-4±0.2	-2±0.3	–	55647.8	0.00219	9.8102×10^{-21}	55.83202(14)	1.569	PSR	
J1453-6413	14:53:32.665(6)	-64:13:16.00(5)	-16±1	-21.3±0.8	–	55433	0.17949	2.74416×10^{-15}	71.248(2)	1.435	PSR	[6]
J1455-3330	14:55:47.9701(12)	-33:30:46.39(4)	7.91±0.04	-1.9±0.09	0.9±0.2	55293	0.00799	2.43043×10^{-20}	13.5698(3)	0.684	PSR	
J1456-6843	14:55:59.9233(3)	-68:43:39.50(2)	-39.5±0.4	-12.3±0.3	2.2±0.3	57600	0.26338	9.895×10^{-17}	8.639(7)	0.438	PSR	
J1502-6752	15:02:18.615(2)	-67:52:16.759(18)	-6±9	-14±16	–	55421.2	0.02674	3.145×10^{-19}	151.2(18)	7.73	PSR	
J1509+5531	15:09:25.6298(1)	+55:31:32.394(2)	-73.64±0.05	-62.65±0.09	0.47±0.03	52275	0.73968	4.99821×10^{-15}	19.6191(3)	2.074	PSR	
J1514-4946	15:14:19.1141(1)	-49:46:15.516(5)	-0.3±3.2	-30±66	–	55520	0.00359	1.8646×10^{-20}	31.05(2)	0.909	PSR	
J1518+0204C	15:18:32.788893(21)	+02:04:47.8153(8)	4.67±0.14	-8.24±0.36	–	52850	0.00248	2.60690×10^{-20}	29.3146(6)	6.616	PSR(GC)	
J1518+4904	15:18:16.799084(16)	+49:04:34.25119(16)	-0.67±0.04	-8.53±0.04	–	52000	0.04093	2.7190×10^{-20}	11.61139(8)	0.964	PSR	
J1527-3931	15:27:58.85(5)	-39:31:35(1)	-33±37	-12±111	–	54217	2.41761	1.90612×10^{-14}	49(3)	1.708	PSR	[7]
J1532+2745	15:32:10.3646(1)	+27:45:49.623(1)	1.542±0.104	18.932±0.111	0.624±0.06	49666	1.12484	7.7956×10^{-16}	14.691(16)	1.321	PSR	
J1534-5334	15:34:08.278(6)	-53:34:19.66(8)	-27±10	-12±10	–	57600	1.36888	1.4288×10^{-15}	24.82(1)	0.81	PSR	
J1537+1155	15:37:09.961730(3)	+11:55:55.43387(6)	1.482±0.007	-25.285±0.012	0.86±0.18	52077	0.03790	2.422494×10^{-18}	11.61944(2)	0.877	PSR	
J1539-5626	15:39:14.06(18)	-56:26:26.3(2)	49±13	-36±16	–	55408	0.24340	4.847844×10^{-15}	175.85(3)	3.536	PSR	
J1543-0620	15:43:30.1373(1)	-06:20:45.332(2)	-16.774±0.044	-0.312±0.131	0.322±0.03	49423	0.70906	8.7954×10^{-16}	18.3774(9)	1.121	PSR	
J1543-0929	15:43:38.8250(1)	+09:29:16.339(2)	-7.61±0.06	-2.87±0.07	0.13±0.02	52275	0.74845	4.3248×10^{-16}	34.9758(16)	5.9*	PSR	
J1543-5149	15:43:44.1498(2)	-51:49:54.685(2)	-4.3±1.4	-4±2	–	55522	0.00206	1.6159×10^{-20}	50.93(14)	1.148	PSR	
J1550-5418	15:50:54.12386(64)	-54:18:24.1141(20)	4.8±0.5	-7.9±0.3	–	54795	2.06983	2.318×10^{-11}	830(50)	6.291	MAG	[6]
J1552-4937	15:52:13.2709(4)	-49:37:49.744(11)	-3±3	-13±8	–	54033	0.00628	1.900×10^{-20}	114.19(8)	3.06	PSR	
J1557-4258	15:57:00.2580(9)	-42:58:12.66(2)	-3±1	10±2	–	48017	0.32919	3.0312×10^{-16}	144.497(2)	8.63	PSR	
J1559-4438	15:59:41.526126(8)	-44:38:45.901778(35)	1.52±0.14	13.15±0.05	0.384±0.08	54100	0.25706	1.01916×10^{-15}	56.1(4)	1.48	PSR	[6]
J1600-3053	16:00:51.903305(13)	-30:53:49.397(61)	-0.97±0.01	-7.04±0.05	0.5±0.06	55885	0.00360	9.50006×10^{-21}	52.3299(2)	2.535	PSR	
J1600-5044	16:00:53.015(2)	-50:44:21.30(4)	-5±2	-13±2	–	57600	0.19260	5.06191×10^{-15}	262.791(4)	5.133	PSR	
J1603-2531	16:03:04.893(3)	-25:31:47.36(7)	-34±4	22±18	–	50719	0.28307	1.593821×10^{-15}	53.763(4)	3.412	PSR	[7]
J1603-7202	16:03:35.67751(4)	-72:02:32.72985(19)	-2.47±0.03	-7.36±0.04	1.1±0.8	54500	0.01484	1.5630×10^{-20}	38.0471(1)	1.129	PSR	[6]
J1604-4909	16:04:22.976(2)	-49:09:57.98(3)	-12±3	9±5	–	57600	0.32742	1.0187×10^{-15}	140.8(5)	3.225	PSR	
J1605+3249	16:05:18.52(3)	+32:49:18.0(3)	-25±16	142±15	–	52111	6.88	–	–	–	XDINS	
J1607-0032	16:07:12.0598(2)	-00:32:41.527(2)	-26.437±0.063	-27.505±0.211	0.934±0.03	46973	0.42182	3.060926×10^{-16}	10.6823(1)	0.68	PSR	
J1614+0737	16:14:40.902(4)	07:37:31.70(1)	-5±12	58±23	–	54192	1.20680	2.35957×10^{-15}	21.3949(3)	1.626	PSR	[7]
J1614-2230	16:14:36.5073(11)	-22:30:31.270(81)	3.6±0.2	-33±1	1.46±0.09	56047	0.00315	9.62419×10^{-21}	34.9179(3)	1.431	PSR	[6]

JName	RA [hms]	DEC [° ′ ″]	μ_{RA} [mas yr⁻¹]	μ_{DEC} [mas yr⁻¹]	parallax [mas]	Pos. Epoch [MJD]	P [s]	$\dot{P}$ [s s⁻¹]	DM [pc cm⁻³]	$d_\odot$ [kpc]	Class(Assoc.)	Ref.
J1622−0315	16:22:59.6285(3)	−03:15:37.328(7)	−15.4 ± 1.7	37.5 ± 5.1	—	56769	0.00385	11.6×10^{-21}	21.4	1.141	PSR	
J1622−6617	16:22:03.6681(4)	−66:17:16.978(6)	−3 ± 2	6 ± 4	—	55253.1	0.02362	5.88×10^{-20}	88.024(9)	4.047	PSR	
J1623−0908	16:23:17.6599(1)	−09:08:48.733(2)	−10.769 ± 0.126	23.509 ± 0.118	0.586 ± 0.10	48715	1.27645	2.580711×10^{-15}	68.183(10)	1.67*	PSR	[5]
J1623−2631	16:23:38.2218(2)	−26:31:53.769(2)	−13.4 ± 1	−25 ± 5	—	48725	0.01108	6.7093×10^{-19}	62.8633(5)	3.651	PSR(GC)	[6]
J1628−3205	16:28:07.02(1)	−32:05:48.7(1)	−6.4 ± 1	−19.81 ± 0.82	1.2 ± 0.56	57205.875	0.00321	—	42.1	1.225	Binary PSR	
J1630+3734	16:30:36.46693(7)	+37:34:42.097(1)	2.4 ± 3.5	−15.9 ± 3.4	—	56136	0.00332	10.77×10^{-21}	14.18009(7)	1.187	PSR	[6]
J1640+2224	16:40:16.7450(22)	+22:24:08.8240(41)	2.08 ± 0.01	−11.34 ± 0.02	0.6 ± 0.4	55366	0.00316	2.81727×10^{-21}	18.462(8)	1.507	PSR	[6]
J1641+8049	16:41:20.830(19)	+80:49:52.910(43)	−11 ± 1	37 ± 3	—	56425	0.00202	8.95×10^{-21}	31.08960(3)	3.035	PSR	[6]
J1643−1224	16:43:38.1618(14)	−12:24:58.68(1)	6.03 ± 0.03	4.1 ± 0.1	0.9 ± 0.3	55330	0.00462	1.84617×10^{-20}	62.4143	0.791	PSR	
J1645−0317	16:45:02.0406(1)	−03:17:57.819(2)	−1.011 ± 0.027	20.523 ± 0.176	0.252 ± 0.23	46515	0.38769	1.780431×10^{-15}	35.7555(8)	1.326	PSR	
J1650−1654	16:50:27.1694(7)	−16:54:42.282(2)	−15.07 ± 0.05	−6.55 ± 0.14	—	50862	1.74955	3.2027×10^{-15}	43.25(15)	1.055	PSR	
J1658−5324	16:58:39.34359(9)	−53:24:07.003(1)	0.2 ± 0.8	4.9 ± 0.23	—	55520	0.00244	1.1154×10^{-20}	30.81(3)	0.88	PSR	
J1703−1846	17:03:51.0915(2)	−18:46:14.845(6)	−0.73 ± 0.08	16.92 ± 0.19	0.35 ± 0.05	49905	0.80434	1.731000×10^{-15}	49.551(15)	1.686	PSR	
J1705−1906	17:05:36.09(17)	−19:06:39(23)	−77 ± 2	−52 ± 17	—	48733	0.29899	4.137877×10^{-15}	22.907(3)	0.747	PSR	
J1708−3426	17:08:57.78(2)	−34:26:44.80(10)	−31 ± 9	123 ± 45	—	54565	0.69212	4.20713×10^{-15}	190.7(3)	4.724	PSR	[7]
J1708−3506	17:08:17.62215(1)	−35:06:22.640(4)	−5.3 ± 0.8	−2 ± 3	—	55132.9	0.00451	1.1421×10^{-20}	146.732(2)	3.32	PSR	
J1709−1640	17:09:26.4413(15)	−16:40:57.73(16)	−2 ± 1	47 ± 10	—	46993	0.65305	6.30985×10^{-15}	24.891(1)	0.562	PSR	
J1709+2313	17:09:05.792(6)	+23:13:27.85(1)	−3.2 ± 0.7	−9.7 ± 0.9	—	51145	0.00463	$0.00000363 \times 10^{-15}$	25.3474(2)	2.179	PSR	
J1710+4923	17:10:04.4410(13)	+49:23:11.478(14)	−50.4 ± 0.2	−44.7 ± 0.2	—	56850	0.00322	1.8211×10^{-20}	7.08493(2)	0.506	PSR	
J1713−0747	17:13:49.53356(24)	+07:47:37.4825(82)	4.924 ± 0.001	−3.913 ± 0.002	0.83 ± 0.02	55391	0.00457	8.52961×10^{-21}	15.917(16)	0.915	PSR	[6]
J1719−1438	17:19:10.07293(5)	−14:38:00.942(4)	1.9 ± 0.4	−11 ± 2	—	55235.5	0.00579	8.043×10^{-21}	36.862(4)	0.341	PSR	
J1720−0212	17:20:57.269(11)	−02:12:23.9(5)	−1 ± 4	−26 ± 5	—	49429	0.47772	8.2798×10^{-17}	66.98(4)	2.357	PSR	
J1721−2457	17:21:05.4979(3)	−24:57:06.17(5)	2 ± 1	−26 ± 14	—	55000	0.00350	5.542×10^{-21}	47.758(19)	1.365	PSR	[6]
J1722−3207	17:22:02.959(2)	−32:07:45.3(2)	−1 ± 5	−40 ± 27	—	57600	0.47716	6.4520×10^{-16}	126.064(8)	2.928	PSR	[6]
J1722−3632	17:22:09.793(2)	−36:32:55.0(1)	−19 ± 7	−110 ± 34	—	53882	0.39918	4.4647×10^{-15}	416.2(2)	3.992	PSR	[7]
J1723−2837	17:23:23.1856(8)	−28:37:57.17(11)	−11.713 ± 0.082	−23.99 ± 0.062	1.077 ± 0.05	55667	0.00186	7.5×10^{-21}	19.688(13)	0.72	Binary PSR	[5]
J1727−2946	17:27:15.09493(17)	−29:46:36.797(17)	0.6 ± 0.9	0 ± 8	—	54723	0.02708	2.4632×10^{-19}	60.74(3)	1.882	PSR	[5]
J1730−2304	17:30:21.66624(8)	−23:04:31.19(2)	20.2 ± 0.3	−1 ± 6	1.7 ± 0.2	54500	0.00812	2.01844×10^{-20}	9.61634	0.512	PSR	[6]
J1731−1847	17:31:17.609823(17)	−18:47:32.666(3)	−1.7 ± 0.3	−6 ± 3	—	55215.1	0.00234	2.5407×10^{-20}	106.4711(6)	4.782	PSR	[6]
J1731−4744	17:31:42.160(5)	−47:44:36.26(6)	73 ± 15	−132 ± 14	—	53000	0.82983	1.63626×10^{-13}	123.056(4)	5.54	PSR	[6]
J1732−5049	17:32:47.766731(19)	−50:49:00.1917(4)	−0.37 ± 0.08	−9.9 ± 0.2	—	54500	0.00531	1.41951×10^{-20}	56.8365	1.875	PSR	[6]
J1733−3716	17:33:26.760(2)	−37:16:55.19(10)	4 ± 9	63 ± 34	—	55359	0.33760	1.504396×10^{-14}	153.18(8)	3.15	PSR	
J1735−0724	17:35:04.9730(1)	−07:24:52.130(1)	0.82 ± 0.06	20.63 ± 0.06	0.15 ± 0.04	49887	0.41933	1.214735×10^{-15}	73.512(4)	0.213	PSR	[6]
J1738+0333	17:38:53.9680(25)	+03:33:10.883(10)	7.12 ± 0.04	5 ± 0.1	0.68 ± 0.05	56258	0.00585	2.411992×10^{-20}	33.77312(4)	1.506	PSR	[6]
J1739−2903	17:39:34.289(1)	−29:03:02.30(2)	10 ± 4	84 ± 41	—	54598	0.32289	7.87490×10^{-15}	138.55(2)	2.911	PSR	[7]
J1740+1311	17:40:07.3455(12)	+13:11:56.69(3)	−22 ± 2	−20 ± 2	—	48262	0.80305	1.451320×10^{-15}	48.66823(43)	4.176	PSR	[6]
J1741−0840	17:41:22.5629(1)	−08:40:31.711(1)	0.41 ± 0.1	6.9 ± 0.09	0.28 ± 0.05	48714	2.04308	2.27471×10^{-15}	74.90(5)	0.222	PSR	[7]
J1741+1351	17:41:31.14473(1)	+13:51:44.116(21)	−8.96 ± 0.03	−7.57 ± 0.04	0.5 ± 0.2	56209	0.00375	3.021648×10^{-20}	24.19871(14)	1.362	PSR	[6]

JName	RA [hms]	DEC [° ′ ″]	μ_{RA} [mas yr⁻¹]	μ_{DEC} [mas yr⁻¹]	parallax [mas]	Pos. Epoch [MJD]	P [s]	$\dot{P}$ [s s⁻¹]	DM [pc cm⁻³]	$d_\odot$ [kpc]	Class(Assoc.)	Ref.
J1741-2054	17:41:57.28(3)	-20:54:11.8(3)	-63±12	-89±9	—	54933	0.41370	1.6983×10^{-14}	4.7(1)	0.273	PSR	
J1741-3927	17:41:18.079(5)	-39:27:38.2(2)	20±15	-6±59	—	57600	0.51221	1.7122×10^{-15}	158.5(6)	4.624	PSR	
J1743-3150	17:43:36.685625(1)	-31:50:22(83)	-11±17	-136±131	—	50241	2.41465	1.20781×10^{-13}	193.05(7)	3.332	PSR	[7]
J1744-1134	17:44:29.405789(11)	-11:34:54.68126(8)	18.797±0.004	-9.41±0.02	2.44±0.05	54500	0.00407	8.93471×10^{-21}	3.13695	0.148	Gamma-/X-ray PSR	[6]
J1744-7619	17:44:00.488(2)	-76:19:14.710(9)	-21±3	-7±3	—	55716	0.00469	9.679×10^{-21}	—	—	Gamma-/X-ray PSR	
J1745-0952	17:45:09.1348(2)	-09:52:39.682(15)	-21.2±1.1	11±5	—	53200	0.01938	9.246×10^{-20}	64.474(14)	0.226	PSR	
J1745-1017	17:45:33.8371(7)	+10:17:52.523(2)	6±1	-5±1	—	55400	0.00265	2.729×10^{-21}	23.970(2)	1.214	PSR	
J1745-2900	17:45:40.1662(8)	-29:00:29.896(1)	2.45±0.32	5.89±0.11	—	56899	3.76373	1.756×10^{-11}	1778(3)	8.289	MAG	
J1745-3040	17:45:56.316(12)	-30:40:22.9(11)	15±1	47±7	—	55276	0.36743	1.067181×10^{-14}	88.373(4)	2.343	PSR	
J1747-4036	17:47:48.7165(17)	-40:36:54.790(82)	-0.2±0.8	-5±1	0.4±0.7	56676	0.00165	1.3142×10^{-20}	152.9595(5)	7.151	PSR	[6]
J1751-2857	17:51:32.693197(17)	-28:57:46.520(3)	-7.4±0.1	-4±1	—	55000	0.00391	1.121×10^{-20}	42.84(3)	1.089	PSR	[6]
J1751-4657	17:51:42.196(2)	-46:57:26.53(5)	-13±5	-59±7	—	57600	0.74235	1.29787×10^{-15}	20.4(4)	0.739	PSR	
J1752-2806	17:52:58.6896(17)	-28:06:37.3(3)	-4±6	-5±5	—	46483	0.56256	8.12913×10^{-15}	50.372(8)	1.335	PSR	
J1754+5201	17:54:22.9068(1)	+52:01:12.244(1)	-3.95±0.05	1.11±0.07	0.16±0.03	49666	2.39140	1.5641×10^{-15}	35.0096(64)	4.173	PSR	
J1756-2251	17:56:46.633812(15)	-22:51:59.35(2)	-2.42±0.08	0±20	1.05±0.55	53563	0.02846	1.017502×10^{-18}	121.196(5)	2.807	PSR	
J1801-1417	18:01:51.073331(19)	-14:17:34.526(2)	-10.89±0.12	-3±1	—	55000	0.00363	0.530×10^{-20}	57.26(4)	1.105	PSR	
J1801-2451	18:01:00.016(8)	-24:51:27.5(2)	-11±9	-1±15	—	53348	0.12492	1.279057×10^{-13}	291.55(5)	3.803	PSR	
J1801-3210	18:01:25.8896(2)	-32:10:53.714(17)	-8±2	-11±10	—	55001.9	0.00745	-4.4×10^{-23}	177.713(4)	6.111	PSR	
J1802-2124	18:02:05.335576(9)	-21:24:03.649(3)	-1.1±0.1	-4±4	1.2±0.7	53453	0.01265	7.256×10^{-20}	149.6258(6)	3.026	PSR	[6]
J1803-2137	18:03:51.4105(1)	-21:37:07.351(10)	11.6±1.8	14.8±2.3	—	51544	0.13367	$1.34359375 \times 10^{-13}$	233.99(5)	3.422	PSR	
J1804-2717	18:04:21.133087(19)	-27:17:31.335(4)	2.6±0.2	-18±2	—	55000	0.00934	4.085×10^{-20}	24.674(5)	0.805	PSR	[6]
J1807-0847	18:07:38.0274(4)	-08:47:43.27(2)	-5±4	1±4	—	48244	0.16373	2.87993×10^{-17}	112.3802(11)	2.7	PSR	
J1808-2024	18:08:39.337(4)	-20:24:39.85(6)	-4.5±1.4	-6.9±2	—	53254	7.55592	5.49×10^{-10}	—	13*	MAG	
J1809-1917	18:09:43.136(2)	-19:17:38.1(5)	-19±6	50±90	—	55366	0.08276	2.553016×10^{-14}	197.1(4)	3.268	PSR	
J1809-1943	18:09:51.08696(3)	-19:43:51.9315(4)	-6.6±0.06	-11.7±1	—	55444	5.54074	2.8281×10^{-12}	178(5)	3.153	MAG	
J1809-2332	18:09:50.25(3)	-23:32:22.68(1)	12±8	-24±6	1.05±0.69	55770	0.14679	3.44207×10^{-14}	—	0.88*	Gamma-/X-ray PSR	
J1810+1744	18:10:37.28038(1)	17:44:37.4719(10)	6.4±1.7	-7.2±2	—	57205.875	0.00166	—	39.7	2.361	Binary PSR	[5]
J1810-2005	18:10:58.9919(6)	-20:05:08.27(14)	0±2	17±37	—	53200	0.03282	1.4742×10^{-19}	241.0(3)	3.514	PSR	
J1811-2405	18:11:19.85315(2)	-24:05:18.365(11)	0.65±0.14	-9.1±5.2	0.2±0.2	55208.5	0.00266	1.37756×10^{-20}	60.6005(17)	1.83	PSR	
J1813-2621	18:13:40.59165(1)	-26:21:57.055(18)	-7.3±0.9	-22±16	—	54058	0.00443	1.2466×10^{-20}	112.524(9)	3.013	PSR	
J1816-4510	18:16:35.93436(7)	+45:10:33.8618(8)	-0.17±0.29	-4.42±0.33	0.22±0.15	57205.875	0.00319	4.310×10^{-20}	38.8874(4)	4.356	Binary PSR	[5]
J1817-3618	18:17:05.857(9)	-36:18:04.6(3)	19±5	-16±17	—	57600	0.38702	2.02298×10^{-15}	94.3(5)	4.399	PSR	
J1820-0427	18:20:52.5934(1)	-04:27:37.712(2)	-7.31±0.07	15.89±0.08	0.35±0.05	57600	0.59808	6.32528×10^{-15}	84.435(17)	2.918	PSR	
J1820-1346	18:20:19.743(9)	13:46:15.0(7)	-17±16	21±95	—	54607	0.92146	4.49522×10^{-15}	776.7(17)	5.902	PSR	
J1823+0550	18:23:30.9731(11)	+05:50:24.31(4)	5±11	-2±4	—	48713	0.75291	2.26734×10^{-16}	66.775(3)	3.088	PSR	
J1823-3021A	18:23:40.48701(36)	-30:21:40.127(44)	0.31±0.24	-8.2±1.7	—	55049	0.00544	3.3750×10^{-18}	86.8797(56)	3.145	PSR(GC)	[7]
J1823-3106	18:23:46.78(67)	-31:06:50.0(62)	18±4	27±18	—	50093	0.28405	2.92704×10^{-15}	50.245(4)	1.586	PSR	
J1824-1945	18:24:00.56(18)	-19:46:03.36(32)	12.2±1	10±14	—	55291	0.18934	5.23563×10^{-15}	224.38(6)	5.599	PSR	

JName	RA [hms]	DEC [° ′ ″]	μ_{RA} [mas yr^{-1}]	μ_{DEC} [mas yr^{-1}]	parallax [mas]	Pos. Epoch [MJD]	P [s]	$\dot{P}$ [s s^{-1}]	DM [pc cm^{-3}]	$d_\odot$ [kpc]	Class(Assoc.)	Ref.
J1824-2452A	18:24:32.00788(3)	-24:52:10.834(8)	-0.2±0.2	-6±4	—	54500	0.00305	1.618836×10^{-18}	119.892	3.701	PSR(GC)	[6]
J1825-0935	18:25:30.594(2)	-09:35:21.2(2)	-13±11	-9±5	—	57600	0.76902	5.235540×10^{-14}	19.3833(9)	0.262	PSR	
J1825-1446	18:25:02.96(17)	-14:46:53.77(7)	11.07±0.1	-27.61±0.1	—	55314	0.27920	2.26696×10^{-14}	352.23(4)	4.44	PSR	
J1826-1334	18:26:13.175(3)	-13:34:46.8(1)	23±2.5	-3.9±3.1	—	52400	0.10149	7.5251738×10^{-14}	231.0(10)	3.606	PSR	
J1829-1751	18:29:43.13(12)	-17:51:04(16)	22±13	-150±130	—	49878	0.30713	5.55174×10^{-15}	217.108(9)	5.937	PSR	
J1832-0827	18:32:37.013(2)	-08:27:03.7(12)	-4±4	20±15	—	55397	0.64733	6.389619×10^{-14}	300.869(10)	4.046	PSR	
J1833-0338	18:33:41.8945(1)	-03:39:04.258(1)	-17.34±0.09	15±0.3	0.4±0.06	57600	0.68673	4.156101×10^{-14}	234.538(13)	5.176	PSR	
J1835-1020	18:35:57.5784(10)	-10:20:04.50(6)	25±5	4±21	—	54369	0.30245	5.91866×10^{-15}	113.7(9)	2.874	PSR	
J1835-1106	18:35:18.41(5)	-11:06:16.1(9)	27±46	56±190	—	55429	0.16592	2.05892×10^{-14}	132.679(3)	3.159	PSR	
J1836-0436	18:36:51.790(2)	-04:36:37.7(1)	-3±2	-1±9	—	49532	0.35424	1.66095×10^{-15}	231.5(3)	4.358	PSR	[7]
J1836-1008	18:36:53.90(23)	-10:08:08(14)	18±65	12±220	—	48880	0.56271	1.179902×10^{-14}	316.98(3)	5.343	PSR	[7]
J1839-0643	18:39:09.788(7)	-06:43:44.5(4)	4±32	-170±111	—	54603	0.44955	3.638×10^{-15}	497.9(16)	4.817	PSR	
J1840-5640	18:40:44.5372(1)	+56:40:54.852(1)	-31.21±0.03	-29.08±0.06	0.66±0.04	48717	1.65286	1.494818×10^{-15}	26.77163(17)	2.189	PSR	
J1841-0912	18:41:55.95(15)	+09:12:07.0(37)	-12±5	38±7	—	48266	0.38132	1.090344×10^{-15}	49.1579(43)	1.664	PSR	
J1843-1113	18:43:41.261917(12)	-11:13:31.0686(7)	-1.91±0.06	-3.2±0.2	0.6±0.4	55000	0.00185	0.9554×10^{-20}	59.964(8)	1.706	PSR	[6]
J1843-1448	18:43:01.3750(3)	-14:48:12.61(3)	10.5±0.19	12±15	—	53934	0.00547	6.209×10^{-21}	114.51(7)	3.472	PSR	
J1844+1454	18:44:54.8946(11)	+14:54:14.12(3)	-9±10	45±6	—	49362	0.37546	1.8723×10^{-15}	41.48555(61)	1.683	PSR	
J1847-0402	18:47:22.842(2)	-04:02:14.4(1)	-1±7	8±19	—	57600	0.59781	5.169112×10^{-14}	141.979(5)	3.419	PSR	
J1848-0123	18:48:23.5895(15)	-01:23:58.33(5)	5±2	-26±6	—	53868	0.65943	5.25182×10^{-15}	159.531(8)	3.533	PSR	
J1853-1303	18:53:57.3183(23)	+13:03:44.0510(41)	-1.63±0.02	-2.96±0.04	0.2±0.3	56553	0.00409	8.7161×10^{-21}	30.56996(8)	1.319	PSR	[6]
J1856-3754	18:56:35.41(16)	-37:54:35.8(2)	326.6±0.5	-61.9±0.4	6.2±0.6	50450	7.05520	2.98×10^{-14}	—	0.16*	XDINS	[3]
J1857+0526	18:57:15.851(2)	+05:26:28.64(2)	19±11	-51±20	—	54619	0.34995	6.9311×10^{-15}	466.4(12)	12.257	PSR	[7]
J1857+0943	18:57:36.39042(12)	+09:43:17.1960(41)	-2.652±0.004	-5.423±0.006	0.9±0.1	55367	0.00536	1.78395×10^{-20}	13.3140(25)	0.767	PSR	[6]
J1900-2600	19:00:47.582(4)	-26:00:43.8(5)	-19.9±0.3	-47.3±0.9	0.5±0.6	48891	0.61221	2.04534×10^{-16}	37.994(5)	1.237	PSR	
J1901+0331	19:01:31.781(2)	+03:31:05.97(7)	-7±2	-45±5	—	54336	0.65545	7.4593×10^{-15}	402.080(12)	6.053	PSR	
J1901+0716	19:01:38.956(5)	07:16:34.4(1)	23±16	-28±33	—	54537	0.64400	2.2872×10^{-15}	252.81(7)	7.237	PSR	[7]
J1901-0906	19:01:53.0087(3)	-09:06:11.146(10)	-7.54±0.04	-18.2±0.2	0.52±0.05	50873	1.78193	1.3829×10^{-15}	72.677(18)	2.892	PSR	[7]
J1902-0556	19:02:42.616(3)	05:56:25.92(6)	-4±5	-1±10	—	54288	0.74658	1.2875167×10^{-14}	177.486(13)	4.198	PSR	[7]
J1902-5105	19:02:02.84821(9)	-51:05:56.9695(8)	-4.8±1.3	-4.4±1.6	—	55520	0.00174	9.1993×10^{-21}	36.25(1)	1.645	PSR	
J1903-0135	19:03:29.986(1)	+01:35:38.53(4)	3±7	-13±14	—	57600	0.72931	4.0250×10^{-15}	245.167(6)	6	PSR	
J1903+0327	19:03:05.793213(1)	+03:27:19.20911(6)	-2.8±0.3	-6.6±0.8	0.3±0.7	55000	0.00215	1.88208×10^{-20}	297.5245(6)	6.122	PSR	[6]
J1903-7051	19:03:38.7935(3)	-70:51:43.461(2)	-8.8±1.6	-16±2	—	56526	0.00360	1.043×10^{-20}	19.66(1)	0.93	PSR	
J1904+0004	19:04:12.7180(15)	+00:04:05.29(4)	8±9	-7±16	—	54377	0.13952	1.1805×10^{-16}	233.61(4)	6.364	PSR	
J1905+0400	19:05:28.273436(16)	+04:00:10.8830(6)	-3.8±0.18	-7.3±0.4	—	53700	0.00378	4.9052×10^{-21}	25.6923(12)	1.064	PSR	
J1905+0709	19:05:53.62(2)	07:09:19.4(6)	-1±34	-12±63	—	49466	0.64804	4.9416×10^{-15}	245.34(10)	4.98	PSR	
J1906+0454	19:06:47.13970(7)	+04:54:09.042(2)	-6±1	-14±3	—	58100	0.00208	2.982×10^{-21}	142.793(2)	4.092	PSR	
J1907+0919	19:07:14.33(1)	+09:19:20.1(2)	0.6±0.2	4.2±0.4	—	51052	5.19835	9.2×10^{-11}	—	—	MAG	[7]
J1907+4002	19:07:34.656(8)	+40:02:05.71(11)	11±4	11±1	—	48713	1.23576	5.40707×10^{-16}	30.966(14)	2.46	PSR	

JName	RA [hms]	DEC [° ′ ″]	μ_{RA} [mas yr⁻¹]	μ_{DEC} [mas yr⁻¹]	parallax [mas]	Pos. Epoch [MJD]	P [s]	$\dot{P}$ [s s⁻¹]	DM [pc cm⁻³]	$d_\odot$ [kpc]	Class(Assoc.)	Ref.
J1909+0007	19:09:35.25(11)	+00:07:58.0(61)	5±3	−32±7	—	48740	1.01695	5.51856×10⁻¹⁵	112.787(6)	4.358	PSR	[6]
J1909-3744	19:09:47.4346749(11)	−37:44:14.46674(5)	−9.513±0.002	−35.777±0.006	0.88±0.01	54500	0.00295	1.402541×10⁻²⁰	10.3932	0.564	PSR	[6]
J1910+1256	19:10:09.7014(35)	+12:56:25.459(82)	0.21±0.03	−7.04±0.05	0.1±0.3	56131	0.00498	9.6807×10⁻²¹	38.0684(5)	1.496	PSR	
J1910-5959A	19:11:42.755623(3)	−59:58:26.9029(3)	−3.08±0.06	−3.97±0.06	—	51920	0.00327	2.94703×10⁻²¹	33.6998(16)	1.685	PSR(GC)	
J1910-5959C	19:11:05.5552(4)	−60:00:59.700(4)	−4.1±1.7	−4.6±2.5	—	51910	0.00527	2.16×10⁻²¹	33.21(4)	1.627	PSR(GC)	[6]
J1911-1114	19:11:49.28233(3)	−11:14:22.481(3)	−13.7±0.2	−9.3±0.9	—	55000	0.00363	1.395×10⁻²⁰	31.02(11)	1.07	PSR	[6]
J1911+1347	19:11:55.20362(12)	+13:47:34.359(21)	2.9±0.03	−3.76±0.05	0.4±0.2	56936	0.00463	1.69287×10⁻²⁰	30.9883(9)	1.365	PSR	
J1912+2104	19:12:43.3391(1)	+21:04:33.926(1)	−11.34±0.06	−5.8±0.1	0.02±0.09	48740	2.23297	1.017751×10⁻¹⁴	88.5961(62)	3.369	PSR	
J1913-0440	19:13:54.1735(16)	−04:40:47.68(7)	7±13	−5±9	—	46634	0.82594	4.06801×10⁻¹⁵	89.385(10)	4.041	PSR	
J1913+1400	19:13:24.3527(1)	+14:00:52.559(1)	−5.28±0.06	−8.94±0.05	0.185±0.03	49910	0.52147	8.03532×10⁻¹⁶	145.052(5)	5.246	PSR	[7]
J1915-1009	19:15:29.991(3)	+10:09:43.37(6)	6±2	−17±4	—	57600	0.40455	1.525768×10⁻¹⁴	241.693(10)	5.99	PSR	
J1915-1606	19:15:27.99942(3)	+16:06:27.3868(5)	−1.23±0.04	−0.83±0.04	—	52984	0.05903	8.6183×10⁻¹⁸	168.77(1)	5.247	PSR	[6]
J1916+0951	19:16:32.3391(9)	09:51:25.93(2)	−1±2	−11±2	—	54587	0.27025	2.518463×10⁻¹⁵	60.953(6)	1.904	PSR	[7]
J1917+1353	19:17:39.7864(1)	+13:53:57.077(1)	−1.25±0.05	3.81±0.06	0.14±0.07	49763	0.19463	7.19782×10⁻¹⁵	94.538(4)	2.94	PSR	
J1918-0642	19:18:48.03272(15)	−06:42:34.9000(82)	−7.15±0.02	−5.94±0.05	0.8±0.1	55330	0.00765	2.569143×10⁻²⁰	26.46(3)	1.024	PSR	
J1918+1444	19:18:23.638(6)	+14:45:06.00(1)	78±49	−262±77	—	49690	1.18102	2.123644×10⁻¹³	27.202(17)	1.304	PSR	
J1919+0021	19:19:50.6715(1)	+00:21:39.722(2)	10.11±0.09	−4.7±0.09	0.17±0.04	49427	1.27226	7.66978×10⁻¹⁵	90.315(16)	4.1	PSR	[6]
J1921+1929	19:21:23.35070(4)	+19:29:22.299(1)	−3.2±0.8	−11±1	—	58137	0.00265	3.8196×10⁻²⁰	64.7632(9)	2.434	PSR	[7]
J1921+2153	19:21:44.815(2)	+21:53:02.25(4)	17±4	32±6	—	48999	1.33730	1.348207×10⁻¹⁵	12.44399(63)	0.807	PSR	
J1923+2515	19:23:22.4926(22)	+25:15:40.5920(61)	−6.69±0.1	−14.7±0.2	−0.3±0.6	56583	0.00379	9.5560×10⁻²¹	18.85515(18)	1.201	PSR	[6]
J1926+0431	19:26:24.471(3)	04:31:31.37(8)	1±9	−14±16	—	54606	1.07408	2.46199×10⁻¹⁵	102.243(19)	4.995	PSR	[7]
J1932+1059	19:32:14.0570(8)	+10:59:33.38(2)	94.09±0.11	42.99±0.16	2.77±0.07	57600	0.22652	1.156826×10⁻¹⁵	3.18321(16)	0.229	PSR	
J1933-6211	19:33:32.42668(5)	−62:11:46.876(1)	−5.54±0.07	10.7±0.2	—	53000	0.00354	3.8707×10⁻²¹	11.520(2)	0.65	PSR	
J1935+1616	19:35:47.8259(1)	+16:16:39.986(2)	1.13±0.13	−16.09±0.15	0.22±0.12	52275	0.35874	6.0025237×10⁻¹⁵	158.521(3)	4.314	PSR	[6]
J1937+1658	19:37:24.90851(8)	+16:58:42.553(2)	−1.8±0.7	−5±1	—	57516	0.00396	1.3202×10⁻²⁰	105.840(1)	3.259	PSR	
J1937+2544	19:37:01.2544(1)	+25:44:13.436(1)	−10.04±0.04	−13.06±0.04	0.32±0.03	49703	0.20098	6.429167×10⁻¹⁶	53.221(5)	2.869	PSR	[6]
J1939+2134	19:39:38.561224(2)	+21:34:59.12570(4)	0.074±0.002	−0.41±0.003	0.19±0.05	55000	0.00156	10.51065×10⁻²⁰	71.0237(13)	2.896	PSR	
J1941-2602	19:41:00.4070(18)	−26:02:05.75(12)	12±2	−10±4	—	50076	0.40286	9.561486×10⁻¹⁶	50.036(3)	3.564	PSR	[7]
J1943-1237	19:43:25.480(4)	−12:37:43.10(3)	−1±8	−46±35	—	54561	0.97243	1.655686×10⁻¹⁵	28.918(15)	1.204	PSR	[6]
J1944-0907	19:44:09.3309(25)	+09:07:23.0060(61)	14.14±0.08	−22.6±0.2	0.1±0.3	56570	0.00519	1.73360×10⁻²⁰	24.3598(5)	1.218	PSR	
J1946-1805	19:46:53.044(6)	+18:05:41.24(11)	1±5	−9±4	—	48790	0.44062	2.4092×10⁻¹⁷	16.1356(73)	1.005	PSR	
J1946-2913	19:46:51.734(8)	−29:13:47.1(5)	19±9	−33±20	—	48736	0.95945	1.488340×10⁻¹⁵	44.309(15)	3.125	PSR	
J1946+3417	19:46:25.131145(2)	+34:17:14.68898(14)	−7.01±0.23	4.51±0.21	—	56819	0.00317	3.118×10⁻²¹	110.2160(3)	6.941	PSR	
J1948+3540	19:48:25.0067(1)	+35:40:11.057(15)	−12.6±0.6	0.7±0.6	—	49449	0.71731	7.06120×10⁻¹⁵	129.36748(78)	7.65	PSR	
J1949+3106	19:49:29.6374604(37)	+31:06:03.80604(6)	−5±5	−5±7	−1±3	56000	0.01314	9.39235×10⁻²⁰	164.12669(24)	7.468	PSR	[6]
J1950+2414	19:50:45.063661(24)	+24:14:56.9639(4)	−2.12±0.18	−3.64±0.19	—	56000	0.00431	1.8805×10⁻²⁰	142.0842(31)	7.269	PSR	
J1952+2630	19:52:36.8401(1)	+26:30:28.073(2)	−6±2	0±3	—	55813	0.02073	4.2721×10⁻¹⁸	315.338(2)	10.033	PSR	
J1952+3252	19:52:58.206(1)	+32:52:40.51(1)	−28.8±0.9	−14.7±0.9	—	52793	0.03953	5.844803×10⁻¹⁵	45.006(19)	3.221	PSR	

JName	RA [hms]	DEC [° ′ ″]	μ_{RA} [mas yr^{-1}]	μ_{DEC} [mas yr^{-1}]	parallax [mas]	Pos. Epoch [MJD]	P [s]	$\dot{P}$ [s s^{-1}]	DM [pc cm^{-3}]	$d_\odot$ [kpc]	Class(Assoc.)	Ref.
J1954+2923	19:54:22.554(2)	+29:23:17.29(4)	25±17	−36±10	–	48719	0.42668	1.711 × 10^{-18}	7.932(7)	0.536	PSR	
J1955+2527	19:55:59.39523(7)	+25:27:03.443(2)	−1.9±0.6	−2.4±0.8	–	54800	0.00487	9.12 × 10^{-21}	209.971(3)	8.177	PSR	
J1955+2908	19:55:27.8753(55)	+29:08:43.436(10)	−1.03±0.06	−4.17±0.08	−1±1	56568	0.00613	2.9737 × 10^{-20}	104.5159(6)	6.305	PSR	[6]
J1955+5059	19:55:18.7637(13)	+50:59:55.292(13)	−23±5	54±5	–	48741	0.51894	1.372366 × 10^{-15}	31.98266(13)	2.186	PSR	
J1955+6708	19:55:38.766(71)	+67:08:15.13(4)	−3±1	10±2	–	56600	0.00857	1.260 × 10^{-20}	57.1478(1)	5.304	PSR	
J1957+2516	19:57:34.6115(50)	+25:16:02.076(5)	−5.7±11	−8.9±1.4	0.69±0.86	57205.875	0.00396	2.744 × 10^{-20}	44.137(3)	2.659	Binary PSR	[5]
J1959+2048	19:59:36.76988(5)	+20:48:15.1222(6)	−16±0.5	−25.8±0.6	–	48196	0.00161	1.68515 × 10^{-20}	29.1168(7)	1.73	PSR	[7]
J2004+3137	20:04:52.279(2)	31:37:09.99(3)	−6±2	−6±3	–	54539	2.11126	7.45504 × 10^{-14}	234.820(8)	7.264	PSR	[7]
J2006-0807	20:06:16.3650(1)	−08:07:02.167(3)	−6.19±0.05	−10.59±0.15	0.42±0.1	49766	0.58087	4.601 × 10^{-17}	32.39(3)	1.706	PSR	[6]
J2010-1323	20:10:45.92111(1)	−13:23:56.083(4)	2.56±0.6	−5.9±0.2	0.2±0.1	56235	0.00522	4.8250 × 10^{-21}	22.177(5)	1.163	PSR	
J2010+3051	20:10:03.9542(1)	+30:51:18.892(1)	−10.3±0.9	−3±1	–	57949	0.00482	4.867 × 10^{-21}	133.756(2)	6.45	PSR	
J2013+3845	20:13:10.367(3)	+38:45:43.31(4)	−32.1±1.7	−25±2	–	49718	0.23019	8.851045 × 10^{-15}	238.217(8)	7.123	PSR	
J2016+1948	20:16:57.44349(6)	+19:47:51.5882(12)	1.28±0.26	2.83±0.34	–	53000	0.06494	4.0063 × 10^{-19}	33.8148(16)	2.162	PSR	[6]
J2017+0603	20:17:22.70520(28)	+06:03:05.564(8)	2.4±0.4	−0.5±0.6	0.5±0.2	56682	0.00290	7.9931 × 10^{-21}	23.92344(9)	1.399	PSR	[6]
J2018+2839	20:18:03.8333(9)	+28:39:54.212(16)	−2.6±0.2	−6.2±0.4	1.03±0.1	46384	0.55795	1.481058 × 10^{-16}	14.1977(6)	0.959	PSR	[6]
J2019+2425	20:19:31.94900(3)	+24:25:15.3038(5)	−8.8±0.5	−19.9±0.7	–	48900	0.00393	7.0237 × 10^{-21}	17.203(5)	1.163	PSR	[6]
J2022+2854	20:22:37.0671(7)	+28:54:23.104(11)	−4.4±0.5	−23.6±0.3	0.37±0.12	49692	0.34340	1.894195 × 10^{-15}	24.63109(18)	1.706	PSR	[6]
J2022+5154	20:22:49.8730(12)	+51:54:50.233(11)	−5.23±0.17	11.5±0.3	0.5±0.07	46640	0.52920	3.06337 × 10^{-15}	22.54968(56)	1.418	PSR	
J2029+3744	20:29:23.882(6)	37:44:08.05(8)	10±11	−10±13	–	54150	1.21680	1.231961 × 10^{-14}	190.66(3)	5.771	PSR	[7]
J2032+4127	20:32:13.119(2)	+41:27:24.31(9)	−2.991±0.048	−0.742±0.055	0.693±0.03	55700	0.14325	1.13 × 10^{-14}	114.67(4)	4.624	Binary PSR	
J2033+1734	20:33:27.5118(86)	+17:34:58.470(16)	−5.9±0.4	−9±0.8	–	56945	0.00595	1.1130 × 10^{-20}	25.0809(9)	1.74	PSR	[6]
J2042+0246	20:42:11.00287(5)	+02:46:14.397(2)	15.1±1.2	−14.1±2.5	–	56520	0.00453	14.03 × 10^{-21}	9.2694(2)	0.64	PSR	
J2043+1711	20:43:20.88164(82)	+17:11:28.908(2)	−5.76±0.06	−10.8±0.1	0.93±0.1	56573	0.00238	5.24370 × 10^{-21}	20.757(14)	1.476	PSR	[6]
J2045-3633	20:45:01.50504(12)	+36:33:01.4033(8)	−2.1±1.4	−2.3±0.8	–	56855	0.03168	5.883 × 10^{-19}	129.5477(17)	5.629	PSR	
J2046-0421	20:46:00.1730(1)	−04:21:26.256(2)	10.76±0.04	−4.26±0.22	0.16±0.03	48739	1.54694	1.471483 × 10^{-15}	35.799(10)	3.267	PSR	
J2046+1540	20:46:39.3373(1)	+15:40:33.558(1)	−10.48±0.06	0.66±0.06	0.31±0.08	48742	1.13829	1.82318 × 10^{-16}	39.81796(49)	3.34	PSR	
J2048-1616	20:48:35.640637(4)	−16:16:44.55350(15)	113.16±0.02	−4.6±0.28	1.05±0.03	54100	1.96157	1.095816 × 10^{-14}	11.456(5)	0.775	PSR	
J2051-0827	20:51:07.519768(18)	−08:27:37.7497(8)	5.63±0.04	2.34±0.28	–	55655	0.00451	1.2733 × 10^{-20}	20.7299(17)	1.469	PSR	
J2053+1718	20:53:49.4809(7)	+17:18:44.662(13)	−1±2.3	6.9±2.8	–	53400	0.11927	2.865 × 10^{-19}	26.980(5)	2.087	PSR	
J2053+4650	20:53:52.62804(7)	+46:50:51.7181(4)	−2.8±0.8	−5.4±0.5	–	56917.5915	0.01259	1.7228 × 10^{-19}	98.0828(6)	3.81	PSR	[7]
J2055-2209	20:55:39.161(2)	22:09:27.09(5)	10±5	−10±9	–	54546	0.81518	1.340139 × 10^{-15}	36.34963(29)	2.998	PSR	
J2055-3630	20:55:31.3521(1)	+36:30:21.469(2)	1.04±0.04	−2.46±0.13	0.17±0.03	52275	0.22151	3.691271 × 10^{-16}	97.4155(56)	5.244	PSR	[7]
J2055-3829	20:55:10.306550(4)	+38:29:30.90571(6)	5.92±0.03	0.79±0.07	–	57900	0.00209	9.996 × 10^{-22}	91.8295(2)	4.588	PSR	
J2108+4441	21:08:20.482(6)	+44:41:48.87(6)	3.5±1.3	1.4±1.4	–	48736	0.41487	8.6210 × 10^{-17}	139.827(11)	4.346	PSR	
J2113+2754	21:13:04.3506(1)	+27:54:01.160(1)	−27.96±0.03	−54.46±0.07	0.7±0.02	48741	1.20285	2.622965 × 10^{-15}	25.11106(18)	1.869	PSR	[6]
J2113+4644	21:13:24.3295(1)	+46:44:08.844(1)	9.5±0.1	8.84±0.08	0.46±0.07	53915	1.01468	7.14612 × 10^{-16}	141.26(9)	4.123	PSR	
J2116+1414	21:16:13.752(3)	+14:14:21.04(7)	8±15	−11±5	–	48471	0.44015	2.89263 × 10^{-16}	56.2044(61)	25	PSR	
J2124-3358	21:24:43.849372(1)	−33:58:44.8500(3)	−14.09±0.03	−50.23±0.07	2.6±0.3	54500	0.00493	2.05705 × 10^{-20}	4.60096	0.36	PSR	[6]

JName	RA [hms]	DEC [° ′ ″]	μ_{RA} [mas yr^{-1}]	μ_{DEC} [mas yr^{-1}]	parallax [mas]	Pos. Epoch [MJD]	P [s]	$\dot{P}$ [s s^{-1}]	DM [pc cm^{-3}]	$d_\odot$ [kpc]	Class(Assoc.)	Ref.
J2129−0429	21:29:45.039(1)	−04:29:05.59(8)	12.34±0.15	10.19±0.14	0.424±0.08	55196	0.00762	−	16.9	1.387	Binary PSR	
J2129+1210A	21:29:58.246512(4)	+12:10:01.2339(1)	−0.54±0.14	−4.33±0.25		55146	0.11066	$-0.0210281 \times 10^{-15}$	67.31(5)	14.2*	PSR(GC)	[6]
J2129+1210B	21:29:58.632(1)	+12:10:00.31(3)	1.7±3.3	−1.9±5.9		50000	0.05613	$0.0095406 \times 10^{-15}$	67.69(12)	14.2*	PSR(GC)	[6]
J2129+1210C	21:30:01.2042(1)	+12:10:38.209(4)	−1.3±0.5	−3.3±1		50000	0.03053	$0.00498789 \times 10^{-15}$	67.13(5)	14.2*	PSR(GC)	
J2129−5721	21:29:22.766966(12)	−57:21:14.21183(12)	9.25±0.02	−9.61±0.03	0.424±0.08	54500	0.00373	2.085328×10^{-20}	31.8509	6.174	PSR	[6]
J2144−3933	21:44:12.060404(45)	−39:33:56.88504(32)	−57.89±0.88	−155.9±0.54	6.05±0.56	54100	8.50983	4.96×10^{-16}	3.35(1)	0.289	PSR	
J2145−0750	21:45:50.4588(1)	−07:50:18.514(4)	−9.58±0.03	−8.87±0.07	1.54±0.1	55000	0.01605	2.9788×10^{-20}	8.99761	0.693	PSR	[6]
J2149+6329	21:49:58.7033(2)	+63:29:44.277(2)	15.81±0.11	11.16±0.19	0.36±0.07	44240	0.38014	1.681×10^{-16}	129.7229(55)	3.885	PSR	
J2150+5247	21:50:37.7499(1)	+52:47:49.556(1)	8.4±0.2	−4.5±0.3	0.03±0.12	49693	0.33221	$1.0105979 \times 10^{-14}$	148.930(4)	3.61	PSR	
J2157+4017	21:57:01.8495(1)	+40:17:45.986(2)	16.13±0.1	4.12±0.12	0.28±0.06	52275	1.52527	3.43256×10^{-15}	71.1239(22)	4.776	PSR	
J2204+2700	22:04:43.609(3)	+27:00:54.69(4)	2±16	−3±3		56805	0.08470	1.31×10^{-19}	35.074(1)	3.196	PSR	
J2212+2933	22:12:23.3444(1)	+29:33:05.411(1)	−5.53±0.08	−11.34±0.11	0.23±0.09	49686	1.00459	4.9517×10^{-16}	74.5213(15)	25	PSR	
J2214+3000	22:14:38.8525(57)	+30:00:38.193(97)	20.7±0.1	−1.6±0.2	1.1±0.2	56610	0.00312	1.47261×10^{-20}	22.5450(25)	1.673	PSR	[6]
J2215+5135	22:15:32.6864(1)	+51:35:36.4063(1)	0.31±0.54	1.88±0.6	0.28±0.36	57205.875	0.00261	3.34×10^{-20}	69.1951(2)	2.773	Binary PSR	[5]
J2219+4754	22:19:48.139(3)	+47:54:53.93(3)	−12±8	−30±6		46599	0.53847	2.765209×10^{-15}	43.4975(5)	2.387	PSR	
J2222−0137	22:22:05.969101(7)	−01:37:15.72441(10)	44.73±0.02	−5.68±0.06	3.742±0.01	55473	0.03282	5.8024×10^{-20}	3.277	0.267	PSR	
J2225+6535	22:25:52.8627(3)	+65:35:36.371(1)	147.23±0.23	126.5±0.1	1.2±0.2	53880	0.68254	9.6606×10^{-15}	36.44362(51)	1.881	PSR	
J2229+2643	22:29:50.8845(1)	+26:43:57.648(13)	−1.7±0.1	−5.9±0.1	0.8±0.6	56937	0.00298	1.516×10^{-21}	22.7271(6)	1.8	PSR	[6]
J2234+0611	22:34:23.073090(2)	+06:11:28.68633(7)	25.3±0.02	9.71±0.05	1.03±0.04	56794.09	0.00358	1.201419×10^{-20}	10.7645(15)	0.854	PSR	
J2235+1506	22:35:43.7001(1)	+15:06:49.060(5)	15±4	10±8		49250	0.05977	1.576×10^{-19}	18.09(5)	1.54	PSR	
J2240+5832	22:40:42.939(4)	+58:32:37.52(4)	−6.1±0.8	−21±0.4		54000	0.13993	1.5358×10^{-14}	263.50(5)	7.275	PSR	
J2241−5236	22:41:42.01850(5)	−52:36:36.2260(3)	17.1±0.1	−3.32±0.05		55044.2	0.00219	6.64×10^{-21}	11.41151(2)	0.963	PSR	[6]
J2248−0101	22:48:26.8859(1)	−01:01:48.085(1)	−10.5±0.07	−17.49±0.19	0.25±0.06	50866	0.47723	6.5935×10^{-16}	29.05(3)	25	MAG	
J2301+5852	23:01:08.29(33)	+58:52:44.45(60)	−6.4±0.6	−2.3±0.6		54852	6.97907	4.7123×10^{-13}		3.3*	MAG	
J2302+4442	23:02:46.9783(78)	+44:42:22.085(95)	−5±4	−7±5	1±1	56675	0.00519	1.378×10^{-20}	13.7279(3)	0.859	PSR	[6]
J2305+3100	23:05:58.3212(1)	+31:00:01.281(1)	−3.7±0.04	−15.63±0.11	0.23±0.03	48714	1.57589	2.89250×10^{-15}	49.5845(12)	4.35*	PSR	
J2308+5547	23:08:13.822(5)	+55:47:36.03(4)	−15±8	0±27		48717	0.47507	1.99491×10^{-16}	46.53905(37)	2.075	PSR	
J2313+4253	23:13:08.6209(1)	+42:53:13.043(2)	24.15±0.1	5.95±0.13	0.93±0.07	52275	0.34943	1.123699×10^{-16}	17.27693(33)	1.108	PSR	[6]
J2317+1439	23:17:09.2364(1)	+14:39:31.265(1)	−1.33±0.3	3.51±0.5	0.1±0.1	54977	0.00345	2.43012×10^{-21}	21.877(9)	2.16	PSR	[6]
J2317+2149	23:17:57.8419(1)	+21:49:48.019(1)	8.49±0.07	0.19±0.14	0.51±0.05	48716	1.44465	1.047047×10^{-15}	20.86959(31)	1.799	PSR	[6]
J2321+6024	23:21:55.20(73)	+60:24:31.0(55)	−17±22	−7±19		49303	2.25649	7.03689×10^{-15}	94.591(18)	2.513	PSR	
J2322+2057	23:22:22.33516(7)	+20:57:02.6772(14)	−18.5±0.3	−15.3±0.5		55000	0.00481	0.9661×10^{-20}	13.36(4)	1.009	PSR	[6]
J2322−2650	23:22:34.64004(3)	−26:50:58.3171(6)	−2.4±0.2	−8.3±0.4		56152	0.00346	5.834×10^{-22}	6.149(2)	0.76	PSR	
J2325+6316	23:25:13.3196(2)	+63:16:52.362(1)	−5.92±0.08	−2.05±0.19	0.01±0.05	48309	1.43631	2.82546×10^{-15}	197.37(5)	4.855	PSR	
J2330−2005	23:30:26.885(7)	−20:05:29.63(17)	74.7±1.9	5±3		49878	1.64362	4.63038×10^{-15}	8.456(2)	0.861	PSR	
J2337+6151	23:37:05.762(1)	+61:51:01.53(7)	−1±18	−15±16		53100	0.49537	1.934498×10^{-13}	58.410(15)	2.079	PSR	
J2339−0533	23:39:38.74(1)	−05:33:05.32(3)	4.15±0.48	−10.31±0.3	0.75±0.25	55100	0.00288	1.4102×10^{-20}	8.717(4)	0.751	Binary PSR	
J2346−0609	23:46:50.4978(1)	−06:09:59.899(2)	37.38±0.03	−20.21±0.09	0.27±0.03	51021	1.18146	1.36305×10^{-15}	22.504(19)	3.7*	PSR	

JName	RA [hms]	DEC [° ′ ″]	μ_{RA} [mas yr^{-1}]	μ_{DEC} [mas yr^{-1}]	parallax [mas]	Pos. Epoch [MJD]	P [s]	$\dot{P}$ [s s^{-1}]	DM [pc cm^{-3}]	$d_\odot$ [kpc]	Class(Assoc.)	Ref.
J2354+6155	23:54:04.7830(1)	+61:55:46.845(1)	22.76 ± 0.05	4.897 ± 0.025	0.41 ± 0.04	49405	0.94478	1.626279×10^{-14}	94.662(6)	2.399	PSR	

Appendix C

Parametrisation of the magnetic-field evolution

As outlined in Section 6.2.2, a key ingredient for the magneto-rotational evolution of radio pulsars is a realistic prescription for the evolution of the dipolar magnetic-field strength, B, up to neutron-star ages of 10^8 yr. While earlier population-synthesis studies have typically either neglected magnetic-field decay entirely, or relied on simplified descriptions invoking decaying exponentials or power laws, we choose a different approach and take advantage of recent progress in modelling the magneto-thermal evolution of neutron-star crusts. In particular, we use a set of five magneto-thermal simulations (Viganò et al., 2021) to fit the early-time magnetic-field evolution which is driven by the combined action of the Hall effect and Ohmic dissipation (see, e.g., Pons and Viganò, 2019, and Section 1.6 for details on these mechanisms).

All five curves, shown as solid lines in Figure 6.1, were simulated with realistic assumption on relevant physics. In particular, the stellar structure and composition are based on the equation of state SLy4 (Douchin and Haensel, 2001) for a neutron star of mass $1.4\,\mathrm{M}_\odot$, resulting in a radius of 11.74 km. The impurity parameter at the highest densities in the inner crust is set to 100 (Pons, Viganò, and Rea, 2013), representing the presence of resistive nuclear pasta phases (see, e.g., Chamel and Haensel, 2008), whereas the impurity profile for other crustal densities matches the results of Carreau, Fantina, and Gulminelli (2020) (see their Fig. 5). Furthermore, the model for the neutron-star envelope is taken from Potekhin, Pons, and Page (2015), while specific parametrisation for the superfluid and superconducting energy gaps (SFB for the crustal neutrons, TToa for the core neutrons and CCDKp for the core protons) were adopted from Ho (2015).

What varies between the different simulations is the initial poloidal magnetic-field strength, B, taking the values $10^{12}, 10^{13}, 10^{14}, 10^{15}, 5 \times 10^{15}$ G, respectively. This also implies different toroidal field strengths, which are typically a factor 10 larger than the poloidal Bs. We observe in Figure 6.1 that those runs with larger magnetic fields decay faster. This is a direct result of the Hall effect which depends on B and acts to redistribute the magnetic-field energy to smaller scales, where it subsequently decays due to Ohmic dissipation. For sources with $B \lesssim 10^{12}$ G and coupled thermal evolution, this Hall cascade does not take place and magnetic fields remain pretty much constant on timescales of the order of 10^6 yr.

Above this timescale, however, current magneto-thermal simulations become unreliable because the implementation of relevant microphysics (Potekhin, Pons, and Page, 2015) is unsuited to old neutron stars with temperatures $\lesssim 10^6$ K. In addition, these simulations focus primarily on the crust and do not include a realistic treatment of the highly uncertain dynamics of the neutron-star core, which should become relevant above $\sim 10^6$ yr. As we require a prescription for the field above 10^6 yr for our population synthesis, we develop a simplified parametrisation for the

late-time magnetic-field evolution that encodes the unknown evolution of the stellar core. As highlighted in Equation (6.12), we assume that field changes at late times can be captured by a power law characterised by the index, a_{late}. This choice is physically motivated because several known magnetic-field evolution mechanisms exhibit the same functional form. For example, Hall-like physics are encoded by $a_{\mathrm{late}} = -1$ (Aguilera, Pons, and Miralles, 2008a), while ambipolar diffusion follows a power law with $a_{\mathrm{late}} = -0.5$ (Goldreich and Reisenegger, 1992).

To directly parametrise the behaviour of the magnetic fields across all relevant B ranges and times t, we describe the field evolution with the following broken power laws:

$$B(t) = B_0 \left(1 + \frac{t}{\tau_1}\right)^{a_1} \left(1 + \frac{t}{\tau_2}\right)^{a_2 - a_1} \left(1 + \frac{t}{\tau_{\mathrm{late}}}\right)^{a_{\mathrm{late}} - a_2} \quad \text{for} \quad \tau_1 < \tau_2 < \tau_{\mathrm{late}},$$

(C.1)

$$B(t) = B_0 \left(1 + \frac{t}{\tau_1}\right)^{a_1} \left(1 + \frac{t}{\tau_{\mathrm{late}}}\right)^{a_{\mathrm{late}} - a_1} \quad \text{for} \quad \tau_1 < \tau_{\mathrm{late}} < \tau_2, \tag{C.2}$$

$$B(t) = B_0 \left(1 + \frac{t}{\tau_{\mathrm{late}}}\right)^{a_{\mathrm{late}}} \quad \text{for} \quad \tau_{\mathrm{late}} < \tau_1 < \tau_2. \tag{C.3}$$

Here, the two timescales $\tau_1 \equiv A_1 B_0^{b_1}$ and $\tau_2 \equiv A_2 B_0^{b_2}$ depend on the initial magnetic field, B_0, while τ_{late} is a constant. The latter together with the free parameters $A_{1,2}, b_{1,2}$ and the power-law indices $a_{1,2}$ can be adjusted to closely fit the numerical simulations. In particular, we choose $\tau_{\mathrm{late}} \approx 2 \times 10^6$ yr, $A_1 = 10^{14}$, $b_1 = -0.8$, $A_2 = 6 \times 10^8$, $b_2 = -0.2$, $a_1 = -0.13$, and $a_2 = -3.0$.

For particularly steep power-law indices, a_{late}, the current prescription, in principle, allows the magnetic field to decay to unrealistically small values in contrast with observations of old millisecond pulsars (Lorimer, 2008). To prevent this, we assume that the magnetic field eventually settles at a constant value, B_{late}, for very late times. In line with detected old neutron stars, we randomly sample the logarithm of B_{late} from a normal distribution with a mean $\mu_{\log B, \mathrm{final}} = 8.5$ and a standard deviation $\sigma_{\log B, \mathrm{final}} = 0.5$ as already outlined previously. The result of this magnetic-field prescription for $a_{\mathrm{late}} = -3.0$ is shown as the dashed lines in Figure 6.1.

Appendix D

Coverage calculation

To validate our neural posterior estimates, we follow Cook, Gelman, and Rubin (2006) who demonstrated that for a well-calibrated posterior distribution, the smallest volume that contains the ground truth, θ, for a given sample in a test data set follows a uniform distribution. This, in turn, implies that the cumulative distribution function of these quantiles across the entire test set forms a diagonal line. The graphical representation of this cumulative distribution function is commonly referred to as the *coverage plot* (see Section 2.5.2 and Figure 6.9). Put differently, if we consider a credibility level $1 - \alpha$, we expect the ground truth, θ, to fall into this region for a fraction $1 - \alpha$ of test samples if the coverage is diagonal.

To calculate the corresponding coverage for our posteriors and assess how well they are calibrated, we take advantage of the amortised nature of our approximate posterior. In particular, for each of our 3,600 test samples, we have access to the ground truth, θ, and the corresponding posterior approximation, $q_{F(x,\phi)}(\theta)$, where $F(x, w)$ represents a trained neural network. To determine the coverage, we need to calculate the quantiles for each θ. In our case, where we infer five magneto-rotational parameters and the posterior, $q_{F(x,w)}(\theta)$, is a five-dimensional probability density function (see Equation (6.27)), we obtain corresponding quantiles by determining the so-called highest-density regions (HDRs), i.e., those regions covering our sample space for a given probability $1 - \alpha$ that have the smallest possible volume (Hyndman, 1996). To obtain these HDRs for each of our test samples, we first compute the total log-posterior at the ground truth, θ, i.e., $\log q_{F(x,w)}(\theta)$. From each posterior, we subsequently draw samples, θ_s, with $s \in \{1, \ldots, S\}$, for which we also individually compute the log-posterior, i.e., $\log q_{F(x,w)}(\theta_s)$. The HDR for a given test sample with ground truth, θ, is now the percentage of samples, θ_s, which satisfy the condition $\log q_{F(x,w)}(\theta_s) > \log q_{F(x,w)}(\theta)$. To compute the cumulative distribution function (coverage) across our test set, we repeat this process iteratively for all 3,600 test samples to determine, for a given credibility level $1 - \alpha$, the fraction of test samples where the HDR is smaller or equal to $1 - \alpha$.

Deviations from the diagonal are present when posterior estimates are either too wide (conservative) or too narrow (over-confident). In the former case, ground truths would be enclosed within a given HDRs more often than expected for the true posterior, while in the latter scenario the opposite applies. The resulting coverage curves would, thus, lie above and below the diagonal, respectively, highlighting the benefit of the coverage plot in validating our posteriors.

Finally note that for our ensemble approach, we calculate the HDR with the ensemble posterior, $\bar{q}(\theta)$, using the condition $\log \bar{q}(\theta_s) > \log \bar{q}(\theta)$. The remaining steps are identical to those outlined above.

Bibliography

Abbott, B. P. et al. (Oct. 2017a). "Gravitational Waves and Gamma-Rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A". In: *ApJ* 848.2, L13, p. L13. DOI: 10.3847/2041-8213/aa920c. arXiv: 1710.05834 [astro-ph.HE].

Abbott, B. P. et al. (Oct. 2017b). "GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral". In: *Phys. Rev. Lett.* 119.16, 161101, p. 161101. DOI: 10.1103/PhysRevLett.119.161101. arXiv: 1710.05832 [gr-qc].

— (Oct. 2017c). "Multi-messenger Observations of a Binary Neutron Star Merger". In: *ApJ* 848.2, L12, p. L12. DOI: 10.3847/2041-8213/aa91c9. arXiv: 1710.05833 [astro-ph.HE].

Abbott, B. P. et al. (Mar. 2020). "GW190425: Observation of a Compact Binary Coalescence with Total Mass $\sim 3.4\,M_\odot$". In: *ApJ* 892.1, L3, p. L3. DOI: 10.3847/2041-8213/ab75f5. arXiv: 2001.01761 [astro-ph.HE].

Abdo, A. A. et al. (Oct. 2013). "The Second Fermi Large Area Telescope Catalog of Gamma-Ray Pulsars". In: *ApJS* 208.2, 17, p. 17. DOI: 10.1088/0067-0049/208/2/17. arXiv: 1305.4385 [astro-ph.HE].

Agarwal, Devansh et al. (June 2020). "FETCH: A deep-learning based classifier for fast transient classification". In: *MNRAS* 497.2, pp. 1661–1674. DOI: 10.1093/mnras/staa1856. arXiv: 1902.06343 [astro-ph.IM].

Aguilera, D. N., J. A. Pons, and J. A. Miralles (July 2008a). "2D Cooling of magnetized neutron stars". In: *A&A* 486.1, pp. 255–271. DOI: 10.1051/0004-6361:20078786. arXiv: 0710.0854 [astro-ph].

Aguilera, Deborah N., José A. Pons, and Juan A. Miralles (Feb. 2008b). "The Impact of Magnetic Field on the Thermal Evolution of Neutron Stars". In: *ApJ* 673.2, p. L167. DOI: 10.1086/527547. arXiv: 0712.1353 [astro-ph].

Allen, Gabrielle et al. (Feb. 2019). "Deep Learning for Multi-Messenger Astrophysics: A Gateway for Discovery in the Big Data Era". In: *preprint*, arXiv:1902.00522, arXiv:1902.00522. arXiv: 1902.00522 [astro-ph.IM].

Alpar, M. Ali, Askin Ankay, and Efe Yazgan (Aug. 2001). "Pulsar Spin-down by a Fallback Disk and the P-P Diagram". In: *ApJ* 557.1, pp. L61–L65. DOI: 10.1086/323140. arXiv: astro-ph/0104287 [astro-ph].

Alpar, M. Ali, Ş. Çalışkan, and Ü. Ertan (Feb. 2013). "Fallback Disks, Magnetars and Other Neutron Stars". In: *Feeding Compact Objects: Accretion on All Scales*. Ed. by C. M. Zhang et al. Vol. 290, pp. 93–100. DOI: 10.1017/S1743921312019291. arXiv: 1211.4721 [astro-ph.HE].

Alsing, Justin et al. (Sept. 2019). "Fast likelihood-free cosmology with neural density estimators and active learning". In: *MNRAS* 488.3, pp. 4440–4458. DOI: 10.1093/mnras/stz1960. arXiv: 1903.00007 [astro-ph.CO].

Archibald, R. F. et al. (Sept. 2016). "A Magnetar-like Outburst from a High-B Radio Pulsar". In: *ApJ* 829.1, L21, p. L21. DOI: 10.3847/2041-8205/829/1/L21. arXiv: 1608.01007 [astro-ph.HE].

Arzoumanian, Z., D. F. Chernoff, and J. M. Cordes (Mar. 2002). "The Velocity Distribution of Isolated Radio Pulsars". In: *ApJ* 568.1, pp. 289–301. DOI: 10.1086/338805. arXiv: astro-ph/0106159 [astro-ph].

Ashton, Gregory et al. (Apr. 2019). "BILBY: A User-friendly Bayesian Inference Library for Gravitational-wave Astronomy". In: *ApJS* 241.2, 27, p. 27. DOI: 10.3847/1538-4365/ab06fc. arXiv: 1811.02042 [astro-ph.IM].

Astropy Collaboration et al. (Oct. 2013). "Astropy: A community Python package for astronomy". In: *A&A* 558, A33, A33. DOI: 10.1051/0004-6361/201322068. arXiv: 1307.6212 [astro-ph.IM].

Astropy Collaboration et al. (Sept. 2018). "The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package". In: *AJ* 156.3, 123, p. 123. DOI: 10.3847/1538-3881/aabc4f. arXiv: 1801.02634 [astro-ph.IM].

Baade, W. and F. Zwicky (July 1934). "Remarks on Super-Novae and Cosmic Rays". In: *Physical Review* 46.1, pp. 76–77. DOI: 10.1103/PhysRev.46.76.2.

Bachetti, M. et al. (Oct. 2014). "An ultraluminous X-ray source powered by an accreting neutron star". In: *Nature* 514.7521, pp. 202–204. DOI: 10.1038/nature13791. arXiv: 1410.3590 [astro-ph.HE].

Backus, George (May 1956). "The External Electric Field of a Rotating Magnet." In: *ApJ* 123, p. 508. DOI: 10.1086/146187.

Bagnulo, S. and J. D. Landstreet (Nov. 2021). "New insight into the magnetism of degenerate stars from the analysis of a volume-limited sample of white dwarfs". In: *MNRAS* 507.4, pp. 5902–5951. DOI: 10.1093/mnras/stab2046. arXiv: 2106.11109 [astro-ph.SR].

Balakrishnan, Vishnu et al. (July 2021). "Pulsar candidate identification using semi-supervised generative adversarial networks". In: *MNRAS* 505.1, pp. 1180–1194. DOI: 10.1093/mnras/stab1308. arXiv: 2010.07457 [astro-ph.IM].

Ball, Nicholas M. and Robert J. Brunner (Jan. 2010). "Data Mining and Machine Learning in Astronomy". In: *Int. J. Mod. Phys. D* 19.7, pp. 1049–1106. DOI: 10.1142/S0218271810017160. arXiv: 0906.2173 [astro-ph.IM].

Balucinska-Church, Monika and Dan McCammon (Dec. 1992). "Photoelectric Absorption Cross Sections with Variable Abundances". In: *ApJ* 400, p. 699. DOI: 10.1086/172032.

Baron, Dalya (Apr. 2019). "Machine Learning in Astronomy: a practical overview". In: *preprint*, arXiv:1904.07248, arXiv:1904.07248. arXiv: 1904.07248 [astro-ph.IM].

Bates, S. D. et al. (Apr. 2014). "PSRPOPPy: an open-source package for pulsar population simulations". In: *MNRAS* 439.3, pp. 2893–2902. DOI: 10.1093/mnras/stu157. arXiv: 1311.3427 [astro-ph.IM].

Beaumont, Mark A., Wenyang Zhang, and David J. Balding (Dec. 2002). "Approximate Bayesian Computation in Population Genetics". In: *Genetics* 162.4, pp. 2025–2035. DOI: 10.1093/genetics/162.4.2025. URL: https://doi.org/10.1093/genetics/162.4.2025.

Beloborodov, Andrei M. (Aug. 2008). "Polar-Cap Accelerator and Radio Emission from Pulsars". In: *ApJ* 683.1, p. L41. DOI: 10.1086/590079. arXiv: 0710.0920 [astro-ph].

Beniamini, Paz, Zorawar Wadiasingh, and Brian D. Metzger (Aug. 2020). "Periodicity in recurrent fast radio bursts and the origin of ultralong period magnetars". In: *MNRAS* 496.3, pp. 3390–3401. DOI: 10.1093/mnras/staa1783. arXiv: 2003.12509 [astro-ph.HE].

Beniamini, Paz et al. (July 2019). "Formation rates and evolution histories of magnetars". In: *MNRAS* 487.1, pp. 1426–1438. DOI: 10.1093/mnras/stz1391. arXiv: 1903.06718 [astro-ph.HE].

Benli, O. and Ü. Ertan (Apr. 2016). "Long-term evolution of anomalous X-ray pulsars and soft gamma repeaters". In: *MNRAS* 457.4, pp. 4114–4122. DOI: 10.1093/mnras/stw235. arXiv: 1601.07846 [astro-ph.HE].

Berger, Edo (Aug. 2014). "Short-Duration Gamma-Ray Bursts". In: *ARA&A* 52, pp. 43–105. DOI: 10.1146/annurev-astro-081913-035926. arXiv: 1311.2603 [astro-ph.HE].

Bessolaz, N. et al. (Jan. 2008). "Accretion funnels onto weakly magnetized young stars". In: *A&A* 478.1, pp. 155–162. DOI: 10.1051/0004-6361:20078328. arXiv: 0712.2921 [astro-ph].

Bethapudi, S. and S. Desai (Apr. 2018). "Separation of pulsar signals from noise using supervised machine learning algorithms". In: *Astronomy and Computing* 23, 15, p. 15. DOI: 10.1016/j.ascom.2018.02.002. arXiv: 1704.04659 [astro-ph.IM].

Bhardwaj, Uddipta et al. (Aug. 2023). "Sequential simulation-based inference for gravitational wave signals". In: *Phys. Rev. D* 108.4, 042004, p. 042004. DOI: 10.1103/PhysRevD.108.042004. arXiv: 2304.02035 [gr-qc].

Bhattacharya, D. and E. P. J. van den Heuvel (Jan. 1991). "Formation and evolution of binary and millisecond radio pulsars". In: *Phys. Rep.* 203.1-2, pp. 1–124. DOI: 10.1016/0370-1573(91)90064-S.

Binney, James and Scott Tremaine (2008). *Galactic Dynamics: Second Edition*.

Bishop, C.M. (1995). *Neural networks for pattern recognition*. Oxford University Press, USA.

Bisnovatyi-Kogan, G. S. (Jan. 1993). "Asymmetric neutrino emission and formation of rapidly moving pulsars". In: *Astronomical and Astrophysical Transactions* 3.4, pp. 287–294. DOI: 10.1080/10556799308230566. arXiv: astro-ph/9707120 [astro-ph].

Bonjean, V. et al. (Feb. 2019). "Star formation rates and stellar masses from machine learning". In: *A&A* 622, A137, A137. DOI: 10.1051/0004-6361/201833972. arXiv: 1901.01932 [astro-ph.GA].

Bovy, Jo (Feb. 2015). "galpy: A python Library for Galactic Dynamics". In: *ApJS* 216.2, 29, p. 29. DOI: 10.1088/0067-0049/216/2/29. arXiv: 1412.3451 [astro-ph.GA].

Brandenburg, Axel (Sept. 2020). "Hall Cascade with Fractional Magnetic Helicity in Neutron Star Crusts". In: *ApJ* 901.1, 18, p. 18. DOI: 10.3847/1538-4357/abad92. arXiv: 2006.12984 [astro-ph.HE].

Braun, Robert, W. Miller Goss, and Andrew G. Lyne (May 1989). "Three Fields Containing Young Pulsars: The Observable Lifetime of Supernova Remnants". In: *ApJ* 340, p. 355. DOI: 10.1086/167398.

Brisken, W. F. et al. (Dec. 2003). "Proper-Motion Measurements with the VLA. II. Observations of 28 Pulsars". In: *AJ* 126.6, pp. 3090–3098. DOI: 10.1086/379559. arXiv: astro-ph/0309215 [astro-ph].

Brisken, Walter F. et al. (June 2002). "Very Long Baseline Array Measurement of Nine Pulsar Parallaxes". In: *ApJ* 571.2, pp. 906–917. DOI: 10.1086/340098. arXiv: astro-ph/0204105 [astro-ph].

Bucciantini, N. et al. (Jan. 2008). "Relativistic jets and long-duration gamma-ray bursts from the birth of magnetars". In: *MNRAS* 383.1, pp. L25–L29. DOI: 10.1111/j.1745-3933.2007.00403.x. arXiv: 0707.2100 [astro-ph].

Buckley, D. A. H. et al. (Jan. 2017). "Polarimetric evidence of a white dwarf pulsar in the binary system AR Scorpii". In: *Nature Astronomy* 1, 0029, p. 0029. DOI: 10.1038/s41550-016-0029. arXiv: 1612.03185 [astro-ph.HE].

Burges, Christopher J. C. (1998). "A Tutorial on Support Vector Machines for Pattern Recognition." In: *Data Min. Knowl. Discov.* 2.2, pp. 121–167. URL: http://dblp.uni-trier.de/db/journals/datamine/datamine2.html#Burges98.

Burrows, A. and J. M. Lattimer (Jan. 1988). "Convection, Type II supernovae, and the early evolution of neutron stars." In: *Phys. Rep.* 163.1, pp. 51–62. DOI: 10.1016/0370-1573(88)90035-X.

Burrows, Adam (July 1987). "Convection and the Mechanism of Type II Supernovae". In: *ApJ* 318, p. L57. DOI: 10.1086/184937.

Cabero, Miriam, Ashish Mahabal, and Jess McIver (Nov. 2020). "GWSkyNet: A Real-time Classifier for Public Gravitational-wave Candidates". In: *ApJ* 904.1, L9, p. L9. DOI: 10.3847/2041-8213/abc5b5. arXiv: 2010.11829 [gr-qc].

Caiazzo, Ilaria et al. (June 2021). "A highly magnetized and rapidly rotating white dwarf as small as the Moon". In: *Nature* 595.7865, pp. 39–42. DOI: 10.1038/s41586-021-03615-y. arXiv: 2107.08458 [astro-ph.SR].

Caleb, Manisha et al. (May 2022). "Discovery of a radio-emitting neutron star with an ultra-long spin period of 76 s". In: *Nature Astronomy*. DOI: 10.1038/s41550-022-01688-x. arXiv: 2206.01346 [astro-ph.HE].

Camilo, Fernando et al. (Aug. 2006). "Transient pulsed radio emission from a magnetar". In: *Nature* 442.7105, pp. 892–895. DOI: 10.1038/nature04986. arXiv: astro-ph/0605429 [astro-ph].

Cannizzo, John K., Hyung Mok Lee, and Jeremy Goodman (Mar. 1990). "The Disk Accretion of a Tidally Disrupted Star onto a Massive Black Hole". In: *ApJ* 351, p. 38. DOI: 10.1086/168442.

Canuto, V., J. Lodenquai, and M. Ruderman (May 1971). "Thomson Scattering in a Strong Magnetic Field". In: *Phys. Rev. D* 3.10, pp. 2303–2308. DOI: 10.1103/PhysRevD.3.2303.

Capano, Collin D. et al. (Mar. 2020). "Stringent constraints on neutron-star radii from multimessenger observations and nuclear theory". In: *Nature Astronomy* 4, pp. 625–632. DOI: 10.1038/s41550-020-1014-6. arXiv: 1908.10352 [astro-ph.HE].

Carlberg, R. G. and K. A. Innanen (Sept. 1987). "Galactic Chaos and the Circular Velocity at the Sun". In: *AJ* 94, p. 666. DOI: 10.1086/114503.

Carreau, T., A. F. Fantina, and F. Gulminelli (Aug. 2020). "Inner crust of a neutron star at the point of crystallization in a multicomponent approach". In: *A&A* 640, A77, A77. DOI: 10.1051/0004-6361/202038347.

Cerutti, Benoît and Andrei M. Beloborodov (July 2017). "Electrodynamics of Pulsar Magnetospheres". In: *Space Sci. Rev.* 207.1-4, pp. 111–136. DOI: 10.1007/s11214-016-0315-7. arXiv: 1611.04331 [astro-ph.HE].

Chadwick, J. (Feb. 1932). "Possible Existence of a Neutron". In: *Nature* 129.3252, p. 312. DOI: 10.1038/129312a0.

Chamel, N. (Sept. 2017). "Superfluidity and Superconductivity in Neutron Stars". In: *Journal of Astrophysics and Astronomy* 38.3, 43, p. 43. DOI: 10.1007/s12036-017-9470-9. arXiv: 1709.07288 [astro-ph.HE].

Chamel, Nicolas and Pawel Haensel (Dec. 2008). "Physics of Neutron Star Crusts". In: *Living Reviews in Relativity* 11.1, 10, p. 10. DOI: 10.12942/lrr-2008-10. arXiv: 0812.3955 [astro-ph].

Chandrasekhar, S. (Mar. 1931). "The highly collapsed configurations of a stellar mass". In: *MNRAS* 91, pp. 456–466. DOI: 10.1093/mnras/91.5.456.

Chatterjee, Pinaki, Lars Hernquist, and Ramesh Narayan (May 2000). "An Accretion Model for Anomalous X-Ray Pulsars". In: *ApJ* 534.1, pp. 373–379. DOI: 10.1086/308748. arXiv: astro-ph/9912137 [astro-ph].

Chatterjee, S. et al. (Sept. 2005). "Getting Its Kicks: A VLBA Parallax for the Hyperfast Pulsar B1508+55". In: *ApJ* 630.1, pp. L61–L64. DOI: 10.1086/491701. arXiv: astro-ph/0509031 [astro-ph].

Chen, B. Q. et al. (July 2019). "The Galactic spiral structure as revealed by O- and early B-type stars". In: *MNRAS* 487.1, pp. 1400–1409. DOI: 10.1093/mnras/stz1357. arXiv: 1905.05542 [astro-ph.GA].

Chen, Kaiyou and Malvin Ruderman (Jan. 1993). "Pulsar Death Lines and Death Valley". In: *ApJ* 402, p. 264. DOI: 10.1086/172129.

Cheung, Damon H. T. et al. (Oct. 2022). "Testing the robustness of simulation-based gravitational-wave population inference". In: *Phys. Rev. D* 106.8, 083014, p. 083014. DOI: 10.1103/PhysRevD.106.083014. arXiv: 2112.06707 [astro-ph.IM].

Chevalier, Roger A. (Nov. 1989). "Neutron Star Accretion in a Supernova". In: *ApJ* 346, p. 847. DOI: 10.1086/168066.

— (Feb. 2005). "Young Core-Collapse Supernova Remnants and Their Supernovae". In: *ApJ* 619.2, pp. 839–855. DOI: 10.1086/426584. arXiv: astro-ph/0409013 [astro-ph].

Cieślar, Marek, Tomasz Bulik, and Stefan Osłowski (Mar. 2020). "Markov Chain Monte Carlo population synthesis of single radio pulsars in the Galaxy". In: *MNRAS* 492.3, pp. 4043–4057. DOI: 10.1093/mnras/staa073.

Coenen, Thijs et al. (Oct. 2014). "The LOFAR pilot surveys for pulsars and fast radio transients". In: *A&A* 570, A60, A60. DOI: 10.1051/0004-6361/201424495. arXiv: 1408.0411 [astro-ph.SR].

Coleman, Matthew S. B. and Adam Burrows (Dec. 2022). "Kicks and induced spins of neutron stars at birth". In: *MNRAS* 517.3, pp. 3938–3961. DOI: 10.1093/mnras/stac2573. arXiv: 2209.02711 [astro-ph.HE].

Colpi, Monica, Ulrich Geppert, and Dany Page (Jan. 2000). "Period Clustering of the Anomalous X-Ray Pulsars and Magnetic Field Decay in Magnetars". In: *ApJ* 529.1, pp. L29–L32. DOI: 10.1086/312448. arXiv: astro-ph/9912066 [astro-ph].

Comella, J. M. et al. (Feb. 1969). "Crab Nebula Pulsar NP 0532". In: *Nature* 221.5179, pp. 453–454. DOI: 10.1038/221453a0.

Connor, Liam and Joeri van Leeuwen (Dec. 2018). "Applying Deep Learning to Fast Radio Burst Classification". In: *AJ* 156.6, 256, p. 256. DOI: 10.3847/1538-3881/aae649. arXiv: 1803.03084 [astro-ph.IM].

Contopoulos, Ioannis, Demosthenes Kazanas, and Christian Fendt (Jan. 1999). "The Axisymmetric Pulsar Magnetosphere". In: *ApJ* 511.1, pp. 351–358. DOI: 10.1086/306652. arXiv: astro-ph/9903049 [astro-ph].

Cook, Samantha R., Andrew Gelman, and Donald B Rubin (2006). "Validation of Software for Bayesian Models Using Posterior Quantiles". In: *Journal of Computational and Graphical Statistics* 15.3, pp. 675–692. DOI: 10.1198/106186006X136976. eprint: https://doi.org/10.1198/106186006X136976. URL: https://doi.org/10.1198/106186006X136976.

Cordes, J. M. and T. J. W. Lazio (July 2002). "NE2001.I. A New Model for the Galactic Distribution of Free Electrons and its Fluctuations". In: *preprint*, arXiv:astro–ph/0207156.

Cordes, J. M. and M. A. McLaughlin (Oct. 2003). "Searches for Fast Radio Transients". In: *ApJ* 596.2, pp. 1142–1154. DOI: 10.1086/378231. arXiv: astro-ph/0304364 [astro-ph].

Coti Zelati, Francesco et al. (Feb. 2018). "Systematic study of magnetar outbursts". In: *MNRAS* 474.1, pp. 961–1017. DOI: 10.1093/mnras/stx2679. arXiv: 1710.04671 [astro-ph.HE].

Cranmer, Miles et al. (June 2020). "Discovering Symbolic Models from Deep Learning with Inductive Biases". In: *preprint*, arXiv:2006.11287, arXiv:2006.11287. arXiv: 2006.11287 [cs.LG].

Cui, Xiang-Han et al. (Nov. 2021). "Statistical tests of young radio pulsars with/without supernova remnants: implying two origins of neutron stars". In: *MNRAS* 508.1, pp. 279–286. DOI: 10.1093/mnras/stab2498. arXiv: 2110.00179 [astro-ph.HE].

Cumming, Andrew (July 2002). "Magnetic field evolution in accreting white dwarfs". In: *MNRAS* 333.3, pp. 589–602. DOI: 10.1046/j.1365-8711.2002.05434.x. arXiv: astro-ph/0202079 [astro-ph].

Cumming, Andrew, Phil Arras, and Ellen Zweibel (July 2004). "Magnetic Field Evolution in Neutron Star Crusts Due to the Hall Effect and Ohmic Decay". In: *ApJ* 609.2, pp. 999–1017. DOI: 10.1086/421324. arXiv: astro-ph/0402392 [astro-ph].

D'Aì, A. et al. (Dec. 2016). "Evidence for the magnetar nature of 1E 161348-5055 in RCW 103". In: *MNRAS* 463.3, pp. 2394–2404. DOI: 10.1093/mnras/stw2023. arXiv: 1607.04264 [astro-ph.HE].

Dall'Osso, S., J. Granot, and T. Piran (June 2012). "Magnetic field decay in neutron stars: from soft gamma repeaters to 'weak-field magnetars'". In: *MNRAS* 422.4, pp. 2878–2903. DOI: 10.1111/j.1365-2966.2012.20612.x. arXiv: 1110.2498 [astro-ph.HE].

Dang, S. J. et al. (June 2020). "Results of 12 yr of Pulsar Timing at Nanshan. I." In: *ApJ* 896.2, 140, p. 140. DOI: 10.3847/1538-4357/ab9082. arXiv: 2005.02200 [astro-ph.HE].

Danilenko, A. et al. (Apr. 2020). "XMM-Newton observations of a gamma-ray pulsar J0633+0632: pulsations, cooling and large-scale emission". In: *MNRAS* 493.2, pp. 1874–1887. DOI: 10.1093/mnras/staa287. arXiv: 2001.10968 [astro-ph.HE].

Daugherty, Joseph K. and Alice K. Harding (Oct. 1986). "Compton Scattering in Strong Magnetic Fields". In: *ApJ* 309, p. 362. DOI: 10.1086/164608.

Davidson, Kris and Jeremiah P. Ostriker (Jan. 1973). "Neutron-Star Accretion in a Stellar Wind: Model for a Pulsed X-Ray Source". In: *ApJ* 179, pp. 585–598. DOI: 10.1086/151897.

Dax, Maximilian et al. (Dec. 2021). "Real-Time Gravitational Wave Science with Neural Posterior Estimation". In: *Phys. Rev. Lett.* 127.24, 241103, p. 241103. DOI: 10.1103/PhysRevLett.127.241103. arXiv: 2106.12594 [gr-qc].

De Grandis, Davide et al. (Nov. 2020). "Three-dimensional Modeling of the Magnetothermal Evolution of Neutron Stars: Method and Test Cases". In: *ApJ* 903.1, 40, p. 40. DOI: 10.3847/1538-4357/abb6f9. arXiv: 2009.04331 [astro-ph.HE].

De Grandis, Davide et al. (June 2021). "X-Ray Emission from Isolated Neutron Stars Revisited: 3D Magnetothermal Simulations". In: *ApJ* 914.2, 118, p. 118. DOI: 10.3847/1538-4357/abfdac. arXiv: 2105.00684 [astro-ph.HE].

De Luca, A. et al. (Aug. 2006). "A Long-Period, Violently Variable X-ray Source in a Young Supernova Remnant". In: *Science* 313.5788, pp. 814–817. DOI: 10.1126/science.1129185. arXiv: astro-ph/0607173 [astro-ph].

Dean, Thomas A. et al. (Mar. 2011). "Parameter Estimation for Hidden Markov Models with Intractable Likelihoods". In: *arXiv e-prints*, arXiv:1103.5399, arXiv:1103.5399. DOI: 10.48550/arXiv.1103.5399. arXiv: 1103.5399 [math.ST].

Dehman, Clara et al. (Jan. 2023). "3D code for MAgneto-Thermal evolution in Isolated Neutron Stars, MATINS: the magnetic field formalism". In: *MNRAS* 518.1, pp. 1222–1242. DOI: 10.1093/mnras/stac2761. arXiv: 2209.12920 [astro-ph.HE].

Deistler, Michael, Pedro J Goncalves, and Jakob H Macke (Oct. 2022). "Truncated proposals for scalable and hassle-free simulation-based inference". In: *arXiv e-prints*, arXiv:2210.04815, arXiv:2210.04815. DOI: 10.48550/arXiv.2210.04815. arXiv: 2210.04815 [stat.ML].

Deller, A. T. et al. (Aug. 2009). "Precision Southern Hemisphere VLBI Pulsar Astrometry. II. Measurement of Seven Parallaxes". In: *ApJ* 701.2, pp. 1243–1257. DOI: 10.1088/0004-637X/701/2/1243. arXiv: 0906.3897 [astro-ph.SR].

Deller, A. T. et al. (Apr. 2019). "Microarcsecond VLBI Pulsar Astrometry with PSRπ II. Parallax Distances for 57 Pulsars". In: *ApJ* 875.2, 100, p. 100. DOI: 10.3847/1538-4357/ab11c7. arXiv: 1808.09046 [astro-ph.IM].

Deutsch, Arnim J. (Jan. 1955). "The electromagnetic field of an idealized star in rigid rotation in vacuo". In: *Annales d'Astrophysique* 18, p. 1.

Dewey, R. J. and J. M. Cordes (Oct. 1987). "Monte Carlo Simulations of Radio Pulsars and Their Progenitors". In: *ApJ* 321, p. 780. DOI: 10.1086/165671.

Dicke, R. H. (July 1946). "The Measurement of Thermal Radiation at Microwave Frequencies". In: *Review of Scientific Instruments* 17.7, pp. 268–275. DOI: 10.1063/1.1770483.

Diehl, Roland et al. (Jan. 2006). "Radioactive ^{26}Al from massive stars in the Galaxy". In: *Nature* 439.7072, pp. 45–47. DOI: 10.1038/nature04364. arXiv: astro-ph/0601015 [astro-ph].

Dieleman, Sander, Kyle W. Willett, and Joni Dambre (June 2015). "Rotation-invariant convolutional neural networks for galaxy morphology prediction". In: *MNRAS* 450.2, pp. 1441–1459. DOI: 10.1093/mnras/stv632. arXiv: 1503.07077 [astro-ph.IM].

Douchin, F. and P. Haensel (Dec. 2001). "A unified equation of state of dense matter and neutron star structure". In: *A&A* 380, pp. 151–167. DOI: 10.1051/0004-6361:20011402. arXiv: astro-ph/0111092 [astro-ph].

du Plessis, Louis et al. (Dec. 2019). "Constraining the Emission Geometry and Mass of the White Dwarf Pulsar AR Sco Using the Rotating Vector Model". In: *ApJ* 887.1, 44, p. 44. DOI: 10.3847/1538-4357/ab4e19. arXiv: 1910.07401 [astro-ph.HE].

Duncan, Robert C. and Christopher Thompson (June 1992). "Formation of Very Strongly Magnetized Neutron Stars: Implications for Gamma-Ray Bursts". In: *ApJ* 392, p. L9. DOI: 10.1086/186413.

— (Apr. 1996). "Magnetars". In: *High Velocity Neutron Stars*. Ed. by R. E. Rothschild and R. E. Lingenfelter. Vol. 366. American Institute of Physics Conference Series, pp. 111–117. DOI: 10.1063/1.50235.

Edwards, R. T. et al. (2001). "The Swinburne intermediate-latitude pulsar survey". In: *MNRAS* 326.1, pp. 358–374.

Eisenbeiss, T. et al. (Mar. 2010). "New photometry and astrometry of the isolated neutron star RX J0720.4-3125 using recent VLT/FORS observations". In: *AN* 331.3, p. 243. DOI: 10.1002/asna.200911301. arXiv: 1001.0545 [astro-ph.HE].

Ekşi, K. Y., L. Hernquist, and R. Narayan (Apr. 2005). "Where Are All the Fallback Disks? Constraints on Propeller Systems". In: *ApJ* 623.1, pp. L41–L44. DOI: 10.1086/429915. arXiv: astro-ph/0501551 [astro-ph].

Elsner, R. F. and F. K. Lamb (Aug. 1977). "Accretion by magnetic neutron stars. I. Magnetospheric structure and stability." In: *ApJ* 215, pp. 897–913. DOI: 10.1086/155427.

Ertan, Ü. et al. (Sept. 2009). "On the Evolution of Anomalous X-ray Pulsars and Soft Gamma-ray Repeaters with Fall Back Disks". In: *ApJ* 702.2, pp. 1309–1320. DOI: 10.1088/0004-637X/702/2/1309. arXiv: 0907.3222 [astro-ph.HE].

Ertl, Thomas et al. (Apr. 2016). "Erratum: "Progenitor-explosion Connection and Remnant Birth Masses for Neutrino-driven Supernovae of Iron-core Progenitors" <A href="/abs/2012ApJ...757...69U">(2012, ApJ, 757, 69)</A>". In: *ApJ* 821.1, 69, p. 69. DOI: 10.3847/0004-637X/821/1/69. arXiv: 1602.06327 [astro-ph.HE].

Esposito, P. et al. (June 2020). "A Very Young Radio-loud Magnetar". In: *ApJ* 896.2, L30, p. L30. DOI: 10.3847/2041-8213/ab9742. arXiv: 2004.04083 [astro-ph.HE].

Esposito, Paolo, Nanda Rea, and Gian Luca Israel (Jan. 2021). "Magnetars: A Short Review and Some Sparse Considerations". In: *Astrophysics and Space Science Library*. Ed. by Tomaso M. Belloni, Mariano Méndez, and Chengmin Zhang. Vol. 461. Astrophysics and Space Science Library, pp. 97–142. DOI: 10.1007/978-3-662-62110-3_3. arXiv: 1803.05716 [astro-ph.HE].

Faucher-Giguère, Claude-André and Victoria M. Kaspi (May 2006). "Birth and Evolution of Isolated Radio Pulsars". In: *ApJ* 643.1, pp. 332–355. DOI: 10.1086/501516. arXiv: astro-ph/0512585 [astro-ph].

Feroz, F., M. P. Hobson, and M. Bridges (Oct. 2009). "MULTINEST: an efficient and robust Bayesian inference tool for cosmology and particle physics". In: *MNRAS* 398.4, pp. 1601–1614. DOI: 10.1111/j.1365-2966.2009.14548.x. arXiv: 0809.3437 [astro-ph].

Ferrario, Lilia, Domitilla de Martino, and Boris T. Gänsicke (Oct. 2015). "Magnetic White Dwarfs". In: *Space Sci. Rev.* 191.1-4, pp. 111–169. DOI: 10.1007/s11214-015-0152-0. arXiv: 1504.08072 [astro-ph.SR].

Ferrario, Lilia and D. T. Wickramasinghe (Jan. 2005). "Magnetic fields and rotation in white dwarfs and neutron stars". In: *MNRAS* 356.2, pp. 615–620. DOI: 10.1111/j.1365-2966.2004.08474.x.

Ferrario, Lilia and Dayal Wickramasinghe (Apr. 2006). "Modelling of isolated radio pulsars and magnetars on the fossil field hypothesis". In: *MNRAS* 367.3, pp. 1323–1328. DOI: 10.1111/j.1365-2966.2006.10058.x. arXiv: astro-ph/0601258 [astro-ph].

Ferrario, Lilia, Dayal Wickramasinghe, and Adela Kawka (Sept. 2020). "Magnetic fields in isolated and interacting white dwarfs". In: *Advances in Space Research* 66.5, pp. 1025–1056. DOI: 10.1016/j.asr.2019.11.012. arXiv: 2001.10147 [astro-ph.SR].

Fluke, Christopher J. and Colin Jacobs (Jan. 2020). "Surveying the reach and maturity of machine learning and artificial intelligence in astronomy". In: *WIREs Data Mining and Knowledge Discovery* 10.2, e1349, e1349. DOI: 10.1002/widm.1349. arXiv: 1912.02934 [astro-ph.IM].

Foreman-Mackey, Daniel et al. (Mar. 2013). "emcee: The MCMC Hammer". In: *PASP* 125.925, p. 306. DOI: 10.1086/670067. arXiv: 1202.3665 [astro-ph.IM].

Frazier, David T. et al. (Dec. 2017). "Approximate Bayesian Forecasting". In: *arXiv e-prints*, arXiv:1712.07750, arXiv:1712.07750. DOI: 10.48550/arXiv.1712.07750. arXiv: 1712.07750 [stat.ME].

Fryer, Chris et al. (May 1999). "What Can the Accretion-induced Collapse of White Dwarfs Really Explain?" In: *ApJ* 516.2, pp. 892–899. DOI: 10.1086/307119. arXiv: astro-ph/9812058 [astro-ph].

Fryer, Christopher L. and Alexander Kusenko (Apr. 2006). "Effects of Neutrino-driven Kicks on the Supernova Explosion Mechanism". In: *ApJS* 163.2, pp. 335–343. DOI: 10.1086/500933. arXiv: astro-ph/0512033 [astro-ph].

Fu, Lei and Xiang-Dong Li (Oct. 2013). "Population Synthesis of Young Isolated Neutron Stars: The Effect of Fallback Disk Accretion and Magnetic Field Evolution". In: *ApJ* 775.2, 124, p. 124. DOI: 10.1088/0004-637X/775/2/124. arXiv: 1308.3011 [astro-ph.HE].

Fujimoto, Yuki, Kenji Fukushima, and Koichi Murase (July 2018). "Methodology study of machine learning for the neutron star equation of state". In: *Phys. Rev. D* 98.2, 023019, p. 023019. DOI: 10.1103/PhysRevD.98.023019. arXiv: 1711.06748 [nucl-th].

Fürst, F. et al. (Nov. 2016). "Discovery of Coherent Pulsations from the Ultraluminous X-Ray Source NGC 7793 P13". In: *ApJ* 831.2, L14, p. L14. DOI: 10.3847/2041-8205/831/2/L14. arXiv: 1609.07129 [astro-ph.HE].

Gaensler, B. M. and Simon Johnston (Dec. 1995). "The pulsar/supernova remnant connection". In: *MNRAS* 277.4, pp. 1243–1253. DOI: 10.1093/mnras/277.4.1243.

Gangadhara, R. T. and Y. Gupta (July 2001). "Understanding the Radio Emission Geometry of PSR B0329+54". In: *ApJ* 555.1, pp. 31–39. DOI: 10.1086/321439. arXiv: astro-ph/0103109 [astro-ph].

García, C. R., Diego F. Torres, and Alessandro Patruno (Sept. 2022). "Visualizing the pulsar population using graph theory". In: *MNRAS* 515.3, pp. 3883–3897. DOI: 10.1093/mnras/stac1997. arXiv: 2207.06311 [astro-ph.HE].

García-Berro, Enrique et al. (2012). "DOUBLE DEGENERATE MERGERS AS PROGENITORS OF HIGH-FIELD MAGNETIC WHITE DWARFS". In: *The Astrophysical Journal* 749.1, p. 25. DOI: 10.1088/0004-637X/749/1/25. URL: https://dx.doi.org/10.1088/0004-637X/749/1/25.

Gavriil, F. P. et al. (Mar. 2008). "Magnetar-Like Emission from the Young Pulsar in Kes 75". In: *Science* 319.5871, p. 1802. DOI: 10.1126/science.1153465. arXiv: 0802.1704 [astro-ph].

Geng, Jin-Jun, Bing Zhang, and Yong-Feng Huang (Nov. 2016). "A Model of White Dwarf Pulsar AR Scorpii". In: *ApJ* 831.1, L10, p. L10. DOI: 10.3847/2041-8205/831/1/L10. arXiv: 1609.02508 [astro-ph.HE].

George, Daniel and E. A. Huerta (Mar. 2018). "Deep Learning for real-time gravitational wave detection and parameter estimation: Results with Advanced LIGO data". In: *Physics Letters B* 778, pp. 64–70. DOI: 10.1016/j.physletb.2017.12.053. arXiv: 1711.03121 [gr-qc].

Gerosa, Davide, Geraint Pratten, and Alberto Vecchio (Nov. 2020). "Gravitational-wave selection effects using neural-network classifiers". In: *Phys. Rev. D* 102.10, 103020, p. 103020. DOI: 10.1103/PhysRevD.102.103020. arXiv: 2007.06585 [astro-ph.HE].

Ghosh, P. and F. K. Lamb (Aug. 1979). "Accretion by rotating magnetic neutron stars. II. Radial and vertical structure of the transition zone in disk accretion." In: *ApJ* 232, pp. 259–276. DOI: 10.1086/157285.

Giacobbo, Nicola and Michela Mapelli (Mar. 2020). "Revising Natal Kick Prescriptions in Population Synthesis Simulations". In: *ApJ* 891.2, 141, p. 141. DOI: 10.3847/1538-4357/ab7335. arXiv: 1909.06385 [astro-ph.HE].

Glorot, Xavier and Yoshua Bengio (May 2010). "Understanding the difficulty of training deep feedforward neural networks". In: *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*. Ed. by Yee Whye Teh and Mike Titterington. Vol. 9. Proceedings of Machine Learning Research. Chia Laguna Resort, Sardinia, Italy: PMLR, pp. 249–256. URL: https://proceedings.mlr.press/v9/glorot10a.html.

Glorot, Xavier, Antoine Bordes, and Yoshua Bengio (Apr. 2011). "Deep Sparse Rectifier Neural Networks". In: *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*. Ed. by Geoffrey Gordon, David Dunson, and Miroslav Dudík. Vol. 15. Proceedings of Machine Learning Research. Fort Lauderdale, FL, USA: PMLR, pp. 315–323. URL: https://proceedings.mlr.press/v15/glorot11a.html.

Gold, T. (May 1968). "Rotating Neutron Stars as the Origin of the Pulsating Radio Sources". In: *Nature* 218.5143, pp. 731–732. DOI: 10.1038/218731a0.

Goldreich, Peter and William H. Julian (Aug. 1969). "Pulsar Electrodynamics". In: *ApJ* 157, p. 869. DOI: 10.1086/150119.

Goldreich, Peter and Andreas Reisenegger (Aug. 1992). "Magnetic Field Decay in Isolated Neutron Stars". In: *ApJ* 395, p. 250. DOI: 10.1086/171646.

Goldstein, A. et al. (Oct. 2017). "An Ordinary Short Gamma-Ray Burst with Extraordinary Implications: Fermi-GBM Detection of GRB 170817A". In: *ApJ* 848.2, L14, p. L14. DOI: 10.3847/2041-8213/aa8f41. arXiv: 1710.05446 [astro-ph.HE].

Gonthier, Peter L. et al. (June 2007). "Population statistics study of radio and gamma-ray pulsars in the Galactic plane". In: *Ap&SS* 309.1-4, pp. 245–251. DOI: 10.1007/s10509-007-9432-9. arXiv: astro-ph/0702097 [astro-ph].

Good, D. C. et al. (Dec. 2020). "First discovery of new pulsars and RRATs with CHIME/FRB". In: *preprint*, arXiv:2012.02320, arXiv:2012.02320. arXiv: 2012.02320 [astro-ph.HE].

Goodfellow, Ian, Yoshua Bengio, and Aaron Courville (2016). *Deep Learning*. http://www.deeplearningbook.org. MIT Press.

Gourgouliatos, K. N. and A. Cumming (Feb. 2014a). "Hall effect in neutron star crusts: evolution, endpoint and dependence on initial conditions". In: *MNRAS* 438.2, pp. 1618–1629. DOI: 10.1093/mnras/stt2300. arXiv: 1311.7004 [astro-ph.SR].

Gourgouliatos, Konstantinos N. and Andrew Cumming (May 2014b). "Hall Attractor in Axially Symmetric Magnetic Fields in Neutron Star Crusts". In: *Phys. Rev. Lett.* 112.17, 171101, p. 171101. DOI: 10.1103/PhysRevLett.112.171101. arXiv: 1311.7345 [astro-ph.SR].

Gourgouliatos, Konstantinos N., Toby S. Wood, and Rainer Hollerbach (Apr. 2016). "Magnetic field evolution in magnetar crusts through three-dimensional simulations". In: *Proceedings of the National Academy of Science* 113.15, pp. 3944–3949. DOI: 10.1073/pnas.1522363113. arXiv: 1604.01399 [astro-ph.SR].

Graber, Vanessa et al. (Oct. 2015). "Magnetic field evolution in superconducting neutron stars". In: *MNRAS* 453.1, pp. 671–681. DOI: 10.1093/mnras/stv1648. arXiv: 1505.00124 [astro-ph.SR].

Green, Stephen R., Christine Simpson, and Jonathan Gair (Nov. 2020). "Gravitational-wave parameter estimation with autoregressive neural network flows". In: *Phys. Rev. D* 102.10, 104057, p. 104057. DOI: 10.1103/PhysRevD.102.104057. arXiv: 2002.07656 [astro-ph.IM].

Greenberg, David S., Marcel Nonnenmacher, and Jakob H. Macke (May 2019). "Automatic Posterior Transformation for Likelihood-Free Inference". In: *arXiv e-prints*, arXiv:1905.07488, arXiv:1905.07488. DOI: 10.48550/arXiv.1905.07488. arXiv: 1905.07488 [cs.LG].

Gullón, M. et al. (Nov. 2015). "Population synthesis of isolated neutron stars with magneto-rotational evolution - II. From radio-pulsars to magnetars". In: *MNRAS* 454.1, pp. 615–625. DOI: 10.1093/mnras/stv1644. arXiv: 1507.05452 [astro-ph.HE].

Gullón, Miguel et al. (Sept. 2014). "Population synthesis of isolated neutron stars with magneto-rotational evolution". In: *MNRAS* 443.3, pp. 1891–1899. DOI: 10.1093/mnras/stu1253. arXiv: 1406.6794 [astro-ph.HE].

Gusakov, M. E., E. M. Kantor, and D. D. Ofengeim (Dec. 2020). "Magnetic field evolution time-scales in superconducting neutron stars". In: *MNRAS* 499.3, pp. 4561–4569. DOI: 10.1093/mnras/staa3160. arXiv: 2010.07673 [astro-ph.HE].

Hahn, ChangHoon et al. (Oct. 2023). "SIMBIG: The First Cosmological Constraints from Non-Gaussian and Non-Linear Galaxy Clustering". In: *arXiv e-prints*, arXiv:2310.15246, arXiv:2310.15246. DOI: 10.48550/arXiv.2310.15246. arXiv: 2310.15246 [astro-ph.CO].

Halpern, J. P. and E. V. Gotthelf (Jan. 2010). "Spin-Down Measurement of PSR J1852+0040 in Kesteven 79: Central Compact Objects as Anti-Magnetars". In: *ApJ* 709.1, pp. 436–446. DOI: 10.1088/0004-637X/709/1/436. arXiv: 0911.0093 [astro-ph.HE].

Hamada, T. and E. E. Salpeter (Nov. 1961). "Models for Zero-Temperature Stars." In: *ApJ* 134, p. 683. DOI: 10.1086/147195.

Hameury, Jean-Marie et al. (Aug. 1998). "Accretion disc outbursts: a new version of an old model". In: *MNRAS* 298.4, pp. 1048–1060. DOI: 10.1046/j.1365-8711.1998.01773.x. arXiv: astro-ph/9803242 [astro-ph].

Han, J. L. et al. (June 2021). "The FAST Galactic Plane Pulsar Snapshot survey: I. Project design and pulsar discoveries". In: *Research in Astronomy and Astrophysics* 21.5, 107, p. 107. DOI: 10.1088/1674-4527/21/5/107. arXiv: 2105.08460 [astro-ph.HE].

Hansen, Brad M. S. and E. Sterl Phinney (Nov. 1997). "The pulsar kick velocity distribution". In: *MNRAS* 291.3, pp. 569–577. DOI: 10.1093/mnras/291.3.569. arXiv: astro-ph/9708071 [astro-ph].

Harding, Alice K. (Dec. 2013). "The neutron star zoo". In: *Frontiers of Physics* 8.6, pp. 679–692. DOI: 10.1007/s11467-013-0285-0. arXiv: 1302.0869 [astro-ph.HE].

Haslam, C. G. T. et al. (July 1981). "A 408 MHz all-sky continuum survey. I - Observations at southern declinations and for the North Polar region." In: *A&A* 100, pp. 209–219.

Haslam, C. G. T. et al. (Jan. 1982). "A 408-MHZ All-Sky Continuum Survey. II. The Atlas of Contour Maps". In: *A&AS* 47, p. 1.

He, C., C. Y. Ng, and V. M. Kaspi (May 2013). "The Correlation between Dispersion Measure and X-Ray Column Density from Radio Pulsars". In: *ApJ* 768.1, 64, p. 64. DOI: 10.1088/0004-637X/768/1/64. arXiv: 1303.5170 [astro-ph.HE].

He, Kaiming et al. (Feb. 2015). "Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification". In: *preprint*, arXiv:1502.01852.

Hermans, Joeri, Volodimir Begy, and Gilles Louppe (Mar. 2019). "Likelihood-free MCMC with Amortized Approximate Ratio Estimators". In: *arXiv e-prints*, arXiv:1903.04057, arXiv:1903.04057. DOI: 10.48550/arXiv.1903.04057. arXiv: 1903.04057 [stat.ML].

Hermans, Joeri et al. (Oct. 2021). "A Trust Crisis In Simulation-Based Inference? Your Posterior Approximations Can Be Unfaithful". In: *arXiv e-prints*, arXiv:2110.06581, arXiv:2110.06581. DOI: 10.48550/arXiv.2110.06581. arXiv: 2110.06581 [stat.ML].

Hernquist, Lars (June 1990). "An Analytical Model for Spherical Galaxies and Bulges". In: *ApJ* 356, p. 359. DOI: 10.1086/168845.

Hewish, A. et al. (Feb. 1968). "Observation of a Rapidly Pulsating Radio Source". In: *Nature* 217.5130, pp. 709–713. DOI: 10.1038/217709a0.

Ho, Wynn C. G. (Sept. 2015). "Magnetic field growth in young glitching pulsars with a braking index". In: *MNRAS* 452.1, pp. 845–851. DOI: 10.1093/mnras/stv1339. arXiv: 1506.03933 [astro-ph.SR].

Ho, Wynn C. G. and Nils Andersson (Jan. 2017). "Ejector and propeller spin-down: how might a superluminous supernova millisecond magnetar become the 6.67 h pulsar in RCW 103". In: *MNRAS* 464.1, pp. L65–L69. DOI: 10.1093/mnrasl/slw186. arXiv: 1608.03149 [astro-ph.SR].

Hobbs, G. et al. (Oct. 2004). "Long-term timing observations of 374 pulsars". In: *MNRAS* 353.4, pp. 1311–1344. DOI: 10.1111/j.1365-2966.2004.08157.x.

Hobbs, G. et al. (July 2005). "A statistical study of 233 pulsar proper motions". In: *MNRAS* 360.3, pp. 974–992. DOI: 10.1111/j.1365-2966.2005.09087.x. arXiv: astro-ph/0504584 [astro-ph].

Holberg, J. B. et al. (Nov. 2016). "The 25 parsec local white dwarf population". In: *MNRAS* 462.3, pp. 2295–2318. DOI: 10.1093/mnras/stw1357.

Horowitz, C. J. et al. (Jan. 2015). "Disordered Nuclear Pasta, Magnetic Field Decay, and Crust Cooling in Neutron Stars". In: *Phys. Rev. Lett.* 114.3, 031102, p. 031102. DOI: 10.1103/PhysRevLett.114.031102. arXiv: 1410.2197 [astro-ph.HE].

Hotan, A. W. et al. (Mar. 2021). "Australian square kilometre array pathfinder: I. system description". In: *Publ. Astron. Soc. Australia* 38, e009, e009. DOI: 10.1017/pasa.2021.1. arXiv: 2102.01870 [astro-ph.IM].

Hotelling, Harold (1933). "Analysis of a complex of statistical variables into principal components." In: *Journal of Educational Psychology* 24, pp. 498–520. URL: https://api.semanticscholar.org/CorpusID:144828484.

Hu, Ren-Yu and Yu-Qing Lou (June 2009). "Magnetized massive stars as magnetar progenitors". In: *MNRAS* 396.2, pp. 878–886. DOI: 10.1111/j.1365-2966.2009.14648.x. arXiv: 0902.3111 [astro-ph.HE].

Hui, C. Y. and W. Becker (Oct. 2006). "Probing the proper motion of the central compact object in Puppis-A with the Chandra high resolution camera". In: *A&A* 457.3, pp. L33–L36. DOI: 10.1051/0004-6361:20065941. arXiv: astro-ph/0606750 [astro-ph].

Huppenkothen, Daniela and Matteo Bachetti (Apr. 2022). "Accurate X-ray timing in the presence of systematic biases with simulation-based inference". In: *MNRAS* 511.4, pp. 5689–5708. DOI: 10.1093/mnras/stab3437. arXiv: 2104.03278 [astro-ph.IM].

Hurley, K. et al. (Apr. 2005). "An exceptionally bright flare from SGR 1806-20 and the origins of short-duration γ-ray bursts". In: *Nature* 434.7037, pp. 1098–1103. DOI: 10.1038/nature03519. arXiv: astro-ph/0502329 [astro-ph].

Hurley-Walker, N. et al. (Jan. 2022a). "A radio transient with unusually slow periodic emission". In: *Nature* 601.7894, pp. 526–530. DOI: 10.1038/s41586-021-04272-x.

Hurley-Walker, N. et al. (July 2023). "A long-period radio transient active for three decades". In: *Nature* 619.7970, pp. 487–490. DOI: 10.1038/s41586-023-06202-5.

Hurley-Walker, Natasha et al. (Apr. 2022b). "GaLactic and Extragalactic All-sky Murchison Widefield Array survey eXtended (GLEAM-X) I: Survey Description and Initial Data Release". In: *arXiv e-prints*, arXiv:2204.12762, arXiv:2204.12762. arXiv: 2204.12762 [astro-ph.GA].

Hyndman, Rob J. (1996). "Computing and Graphing Highest Density Regions". In: *The American Statistician* 50.2, pp. 120–126. ISSN: 00031305. URL: http://www.jstor.org/stable/2684423 (visited on 09/11/2023).

Igoshev, Andrei P. (Apr. 2020). "The observed velocity distribution of young pulsars - II. Analysis of complete PSRπ". In: *MNRAS* 494.3, pp. 3663–3674. DOI: 10.1093/mnras/staa958. arXiv: 2002.01367 [astro-ph.HE].

Igoshev, Andrei P. et al. (Mar. 2021). "3D Magnetothermal Simulations of Tangled Crustal Magnetic Field in Central Compact Objects". In: *ApJ* 909.2, 101, p. 101. DOI: 10.3847/1538-4357/abde3e. arXiv: 2101.08292 [astro-ph.HE].

Igoshev, Andrei P. et al. (Aug. 2022). "Initial periods and magnetic fields of neutron stars". In: *MNRAS* 514.3, pp. 4606–4619. DOI: 10.1093/mnras/stac1648. arXiv: 2205.06823 [astro-ph.HE].

Inutsuka, Shu-ichiro and Takayoshi Sano (Aug. 2005). "Self-sustained Ionization and Vanishing Dead Zones in Protoplanetary Disks". In: *ApJ* 628.2, pp. L155–L158. DOI: 10.1086/432796. arXiv: astro-ph/0506131 [astro-ph].

Isern, Jordi et al. (Feb. 2017). "A Common Origin of Magnetism from Planets to White Dwarfs". In: *ApJ* 836.2, L28, p. L28. DOI: 10.3847/2041-8213/aa5eae. arXiv: 1702.01813 [astro-ph.SR].

Israel, Gian Luca et al. (Feb. 2017). "An accreting pulsar with extreme properties drives an ultraluminous x-ray source in NGC 5907". In: *Science* 355.6327, pp. 817–819. DOI: 10.1126/science.aai8635. arXiv: 1609.07375 [astro-ph.HE].

Izvekova, V. A. et al. (Aug. 1981). "Radio Spectra of Pulsars - Part One - Observations of Flux Densities at Meter Wavelengths and Analysis of the Spectra". In: *Ap&SS* 78.1, pp. 45–72. DOI: 10.1007/BF00654022.

Jackson, John David (1998). *Classical Electrodynamics, 3rd Edition*.

Jacoby, B. A. et al. (July 2009). "A Large-Area Survey for Radio Pulsars at High Galactic Latitudes". In: *ApJ* 699.2, pp. 2009–2016. DOI: 10.1088/0004-637X/699/2/2009.

Janka, Hans-Thomas (Mar. 2017). "Neutron Star Kicks by the Gravitational Tug-boat Mechanism in Asymmetric Supernova Explosions: Progenitor and Explosion Dependence". In: *ApJ* 837.1, 84, p. 84. DOI: 10.3847/1538-4357/aa618e. arXiv: 1611.07562 [astro-ph.HE].

Janka, Hans-Thomas, Annop Wongwathanarat, and Michael Kramer (Feb. 2022). "Supernova Fallback as Origin of Neutron Star Spins and Spin-kick Alignment". In: *ApJ* 926.1, 9, p. 9. DOI: 10.3847/1538-4357/ac403c. arXiv: 2104.07493 [astro-ph.HE].

Jankowski, F. et al. (Feb. 2018). "Spectral properties of 441 radio pulsars". In: *MNRAS* 473.4, pp. 4436–4458. DOI: 10.1093/mnras/stx2476. arXiv: 1709.08864 [astro-ph.HE].

Jeffrey, Niall, Justin Alsing, and François Lanusse (Feb. 2021). "Likelihood-free inference with neural compression of DES SV weak lensing map statistics". In: *MNRAS* 501.1, pp. 954–969. DOI: 10.1093/mnras/staa3594. arXiv: 2009.08459 [astro-ph.CO].

Jennings, Ross J. et al. (Sept. 2018). "Binary Pulsar Distances and Velocities from Gaia Data Release 2". In: *ApJ* 864.1, 26, p. 26. DOI: 10.3847/1538-4357/aad084. arXiv: 1806.06076 [astro-ph.SR].

Jimenez Rezende, Danilo and Shakir Mohamed (May 2015). "Variational Inference with Normalizing Flows". In: *arXiv e-prints*, arXiv:1505.05770, arXiv:1505.05770. DOI: 10.48550/arXiv.1505.05770. arXiv: 1505.05770 [stat.ML].

Johnson, J. B. (July 1928). "Thermal Agitation of Electricity in Conductors". In: *Physical Review* 32.1, pp. 97–109. DOI: 10.1103/PhysRev.32.97.

Johnston, S. et al. (Dec. 2008). "Science with ASKAP. The Australian square-kilometre-array pathfinder". In: *Experimental Astronomy* 22.3, pp. 151–273. DOI: 10.1007/s10686-008-9124-7. arXiv: 0810.5187 [astro-ph].

Johnston, Simon and A. Karastergiou (May 2019). "The period-width relationship for radio pulsars revisited". In: *MNRAS* 485.1, pp. 640–647. DOI: 10.1093/mnras/stz400. arXiv: 1902.03996 [astro-ph.HE].

Johnston, Simon et al. (Apr. 1992). "A high-frequency survey of the southern galactic plane for pulsars." In: *MNRAS* 255, pp. 401–411. DOI: 10.1093/mnras/255.3.401.

Johnston, Simon et al. (July 2020). "The Galactic population and properties of young, highly energetic pulsars". In: *MNRAS* 497.2, pp. 1957–1965. DOI: 10.1093/mnras/staa2110. arXiv: 2007.08681 [astro-ph.HE].

Jonas, J. L. (Aug. 2009). "MeerKAT - The South African Array With Composite Dishes and Wide-Band Single Pixel Feeds". In: *IEEE Proceedings* 97.8, pp. 1522–1530. DOI: 10.1109/JPROC.2009.2020713.

Kapoor, R. C. and C. S. Shukre (July 1998). "Kinematics of Pulsar Beams". In: *ApJ* 501.1, pp. 228–241. DOI: 10.1086/305804.

Kasen, Daniel and Lars Bildsten (July 2010). "Supernova Light Curves Powered by Young Magnetars". In: *ApJ* 717.1, pp. 245–249. DOI: 10.1088/0004-637X/717/1/245. arXiv: 0911.0680 [astro-ph.HE].

Kaspi, V. M. et al. (Oct. 2001). "Chandra X-Ray Observations of G11.2-0.3: Implications for Pulsar Ages". In: *ApJ* 560.1, pp. 371–377. DOI: 10.1086/322515. arXiv: astro-ph/0107292 [astro-ph].

Kaspi, Victoria M. (Apr. 2010). "Grand unification of neutron stars". In: *Proceedings of the National Academy of Science* 107.16, pp. 7147–7152. DOI: 10.1073/pnas.1000812107. arXiv: 1005.0876 [astro-ph.HE].

Kaspi, Victoria M. and Andrei M. Beloborodov (Aug. 2017). "Magnetars". In: *ARA&A* 55.1, pp. 261–301. DOI: 10.1146/annurev-astro-081915-023329. arXiv: 1703.00068 [astro-ph.HE].

Kawka, A. and S. Vennes (Dec. 2004). "Ap stars as progenitors of magnetic white dwarfs". In: *The A-Star Puzzle*. Ed. by Juraj Zverko et al. Vol. 224, pp. 879–885. DOI: 10.1017/S1743921305009920.

Keane, E. F. and M. Kramer (Dec. 2008). "On the birthrates of Galactic neutron stars". In: *MNRAS* 391.4, pp. 2009–2016. DOI: 10.1111/j.1365-2966.2008.14045.x. arXiv: 0810.1512 [astro-ph].

Keith, M. J. et al. (Nov. 2010). "The High Time Resolution Universe Pulsar Survey – I. System configuration and initial discoveries". In: *MNRAS* 409.2, pp. 619–627. ISSN: 0035-8711. DOI: 10.1111/j.1365-2966.2010.17325.x. eprint: https://academic.oup.com/mnras/article-pdf/409/2/619/18579028/mnras0409-0619.pdf. URL: https://doi.org/10.1111/j.1365-2966.2010.17325.x.

Kelley, Henry J. (1960). "Gradient Theory of Optimal Flight Paths". In: *ARS Journal* 30.10, pp. 947–954. DOI: 10.2514/8.5282. eprint: https://doi.org/10.2514/8.5282. URL: https://doi.org/10.2514/8.5282.

Kepler, S. O. et al. (Mar. 2013). "Magnetic white dwarf stars in the Sloan Digital Sky Survey". In: *MNRAS* 429.4, pp. 2934–2944. DOI: 10.1093/mnras/sts522. arXiv: 1211.5709 [astro-ph.SR].

Khullar, Gourav et al. (Dec. 2022). "DIGS: deep inference of galaxy spectra with neural posterior estimation". In: *Machine Learning: Science and Technology* 3.4, 04LT04, 04LT04. DOI: 10.1088/2632-2153/ac98f4. arXiv: 2211.09126 [astro-ph.GA].

Kiel, Paul D. and Jarrod R. Hurley (June 2009). "Populating the Galaxy with pulsars - II. Galactic dynamics". In: *MNRAS* 395.4, pp. 2326–2346. DOI: 10.1111/j.1365-2966.2009.14711.x. arXiv: 0903.1905 [astro-ph.GA].

Kiel, Paul D. et al. (July 2008). "Populating the Galaxy with pulsars - I. Stellar and binary evolution". In: *MNRAS* 388.1, pp. 393–415. DOI: 10.1111/j.1365-2966.2008.13402.x. arXiv: 0805.0059 [astro-ph].

Kingma, Diederik P. and Jimmy Ba (Dec. 2014). "Adam: A Method for Stochastic Optimization". In: *preprint*, arXiv:1412.6980, arXiv:1412.6980. arXiv: 1412.6980 [cs.LG].

Kolmogorov, Andrey Nikolaevich (Apr. 1941). "Dissipation of Energy in Locally Isotropic Turbulence". In: *Akademiia Nauk SSSR Doklady* 32, p. 16.

Kramer, M. et al. (Nov. 1994). "Geometrical analysis of average pulsar profiles using multi-component Gaussian FITS at several frequencies. I. Method and analysis." In: *A&AS* 107, pp. 515–526.

Kramer, M. et al. (May 2007). "Polarized radio emission from a magnetar". In: *MNRAS* 377.1, pp. 107–119. DOI: 10.1111/j.1365-2966.2007.11622.x. arXiv: astro-ph/0702365 [astro-ph].

Krishnakumar, M. A. et al. (May 2015). "Scatter Broadening Measurements of 124 Pulsars At 327 Mhz". In: *ApJ* 804.1, 23, p. 23. DOI: 10.1088/0004-637X/804/1/23. arXiv: 1501.05401 [astro-ph.HE].

Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton (2012). "ImageNet Classification with Deep Convolutional Neural Networks". In: *Advances in Neural Information Processing Systems 25*. Ed. by F. Pereira et al. Curran Associates, Inc., pp. 1097–1105. URL: http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf.

Kullback, S. and R. A. Leibler (1951). "On Information and Sufficiency". In: *The Annals of Mathematical Statistics* 22.1, pp. 79 –86. DOI: 10.1214/aoms/1177729694. URL: https://doi.org/10.1214/aoms/1177729694.

Lai, Dong, David F. Chernoff, and James M. Cordes (Mar. 2001). "Pulsar Jets: Implications for Neutron Star Kicks and Initial Spins". In: *ApJ* 549.2, pp. 1111–1118. DOI: 10.1086/319455. arXiv: astro-ph/0007272 [astro-ph].

Lam, Siu Kwan, Antoine Pitrou, and Stanley Seibert (2015). "Numba: A LLVM-Based Python JIT Compiler". In: *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*. LLVM '15. Austin, Texas: Association for Computing

Machinery. ISBN: 9781450340052. URL: https://doi.org/10.1145/2833157.2833162.

Large, M. I., A. E. Vaughan, and B. Y. Mills (Oct. 1968). "A Pulsar Supernova Association?" In: *Nature* 220.5165, pp. 340–341. DOI: 10.1038/220340a0.

Lawson, K. D. et al. (Mar. 1987). "Variations in the spectral index of the galactic radio continuum emission in the northern hemisphere." In: *MNRAS* 225, pp. 307–327. DOI: 10.1093/mnras/225.2.307.

Leahy, D. A., S. Ranasinghe, and M. Gelowitz (May 2020). "Evolutionary Models for 43 Galactic Supernova Remnants with Distances and X-Ray Spectra". In: *ApJS* 248.1, 16, p. 16. DOI: 10.3847/1538-4365/ab8bd9. arXiv: 2003.08998 [astro-ph.HE].

LeCun, Y. et al. (1989). "Backpropagation Applied to Handwritten Zip Code Recognition". In: *Neural Computation* 1.4, pp. 541–551. DOI: 10.1162/neco.1989.1.4.541.

Lemos, Pablo et al. (Mar. 2023). "Robust simulation-based inference in cosmology with Bayesian neural networks". In: *Machine Learning: Science and Technology* 4.1, 01LT01, 01LT01. DOI: 10.1088/2632-2153/acbb53. arXiv: 2207.08435 [astro-ph.CO].

Lentati, L. et al. (May 2016). "From spin noise to systematics: stochastic processes in the first International Pulsar Timing Array data release". In: *MNRAS* 458.2, pp. 2161–2187. DOI: 10.1093/mnras/stw395. arXiv: 1602.05570 [astro-ph.IM].

Levin, L. et al. (Sept. 2013). "The High Time Resolution Universe Pulsar Survey - VIII. The Galactic millisecond pulsar population". In: *MNRAS* 434.2, pp. 1387–1397. DOI: 10.1093/mnras/stt1103. arXiv: 1306.4190 [astro-ph.SR].

Li, Chengdong et al. (Feb. 2019). "Flare and Warp of Galactic Disk with OB Stars from Gaia DR2". In: *ApJ* 871.2, 208, p. 208. DOI: 10.3847/1538-4357/aafa17.

Li, Shao-Ze et al. (Feb. 2021). "The Evolution of a Newborn Millisecond Magnetar with a Propeller-recycling Disk". In: *ApJ* 907.2, 87, p. 87. DOI: 10.3847/1538-4357/abcc70. arXiv: 2011.09762 [astro-ph.HE].

Li, Weidong et al. (Apr. 2011). "Nearby supernova rates from the Lick Observatory Supernova Search - III. The rate-size relation, and the rates as a function of galaxy Hubble type and colour". In: *MNRAS* 412.3, pp. 1473–1507. DOI: 10.1111/j.1365-2966.2011.18162.x. arXiv: 1006.4613 [astro-ph.SR].

Li, Xiang-Dong (Aug. 2006). "On the Nature of Part-Time Radio Pulsars". In: *ApJ* 646.2, pp. L139–L142. DOI: 10.1086/506962. arXiv: astro-ph/0606508 [astro-ph].

— (Sept. 2007). "The Nature of the Compact X-Ray Source in Supernova Remnant RCW 103". In: *ApJ* 666.2, pp. L81–L84. DOI: 10.1086/521791. arXiv: 0708.0060 [astro-ph].

Liebert, James (Nov. 1988). "Searches for Magnetic Fields in White Dwarfs Since Babcock". In: *PASP* 100, p. 1302. DOI: 10.1086/132322.

Lin, Haitao, Xiangru Li, and Ziying Luo (Apr. 2020). "Pulsars detection by machine learning with very few features". In: *MNRAS* 493.2, pp. 1842–1854. DOI: 10.1093/mnras/staa218. arXiv: 2002.08519 [astro-ph.IM].

Lin, Kiyam et al. (Oct. 2023). "A simulation-based inference pipeline for cosmic shear with the Kilo-Degree Survey". In: *MNRAS* 524.4, pp. 6167–6180. DOI: 10.1093/mnras/stad2262. arXiv: 2212.04521 [astro-ph.CO].

Lin, Weili et al. (June 2021). "Supernova Luminosity Powered by Magnetar-Disk System". In: *ApJ* 914.1, L2, p. L2. DOI: 10.3847/2041-8213/ac004a. arXiv: 2105.05512 [astro-ph.HE].

Liu, Bai Sheng and Xiang-Dong Li (Nov. 2015). "On the Evolution of a Fossil Disk around Neutron Stars Originating from Merging WDs". In: *ApJ* 814.1, 75, p. 75. DOI: 10.1088/0004-637X/814/1/75. arXiv: 1510.07404 [astro-ph.HE].

Liu, Jie et al. (Aug. 2020). "New timing measurement results of 16 pulsars". In: *PASJ* 72.4, 70, p. 70. DOI: `10.1093/pasj/psaa058`. arXiv: `2006.02017 [astro-ph.HE]`.

Long, M., M. M. Romanova, and R. V. E. Lovelace (Dec. 2005). "Locking of the Rotation of Disk-Accreting Magnetized Stars". In: *ApJ* 634.2, pp. 1214–1222. DOI: `10.1086/497000`. arXiv: `astro-ph/0510659 [astro-ph]`.

Longair, Malcolm S. (2011). *High Energy Astrophysics*.

Lorimer, D. R. (Jan. 2004). "The Galactic Population and Birth Rate of Radio Pulsars". In: *Young Neutron Stars and Their Environments*. Ed. by Fernando Camilo and Bryan M. Gaensler. Vol. 218, p. 105. DOI: `10.48550/arXiv.astro-ph/0308501`. arXiv: `astro-ph/0308501 [astro-ph]`.

Lorimer, D. R. and M. Kramer (2012). *Handbook of Pulsar Astronomy*. Cambridge University Press.

Lorimer, D. R. et al. (July 1993). "Pulsar statistics : the birthrate and initial spin periods of radio pulsars." In: *MNRAS* 263, pp. 403–415. DOI: `10.1093/mnras/263.2.403`.

Lorimer, D. R. et al. (Oct. 2006). "The Parkes Multibeam Pulsar Survey - VI. Discovery and timing of 142 pulsars and a Galactic population analysis". In: *MNRAS* 372.2, pp. 777–800. DOI: `10.1111/j.1365-2966.2006.10887.x`. arXiv: `astro-ph/0607640 [astro-ph]`.

Lorimer, Duncan R. (Nov. 2008). "Binary and Millisecond Pulsars". In: *Living Reviews in Relativity* 11.1, 8, p. 8. DOI: `10.12942/lrr-2008-8`. arXiv: `0811.0762 [astro-ph]`.

Lueckmann, Jan-Matthis et al. (Nov. 2017). "Flexible statistical inference for mechanistic models of neural dynamics". In: *arXiv e-prints*, arXiv:1711.01861, arXiv:1711.01861. DOI: `10.48550/arXiv.1711.01861`. arXiv: `1711.01861 [stat.ML]`.

Maciesiak, Krzysztof and Janusz Gil (Oct. 2011). "On the pulse-width statistics in radio pulsars - II. Importance of the core profile components". In: *MNRAS* 417.2, pp. 1444–1453. DOI: `10.1111/j.1365-2966.2011.19359.x`. arXiv: `1107.0212 [astro-ph.GA]`.

Maciesiak, Krzysztof, Janusz Gil, and Giorgi Melikidze (Aug. 2012). "On the pulse-width statistics in radio pulsars - III. Importance of the conal profile components". In: *MNRAS* 424.3, pp. 1762–1773. DOI: `10.1111/j.1365-2966.2012.21246.x`. arXiv: `1205.1297 [astro-ph.GA]`.

Maciesiak, Krzysztof, Janusz Gil, and Valério A. R. M. Ribeiro (June 2011). "On the pulse-width statistics in radio pulsars - I. Importance of the interpulse emission". In: *MNRAS* 414.2, pp. 1314–1328. DOI: `10.1111/j.1365-2966.2011.18471.x`. arXiv: `1102.3348 [astro-ph.GA]`.

MacQueen, J. B. (1967). "Some Methods for Classification and Analysis of MultiVariate Observations". In: *Proc. of the fifth Berkeley Symposium on Mathematical Statistics and Probability*. Ed. by L. M. Le Cam and J. Neyman. Vol. 1. University of California Press, pp. 281–297.

Manchester, R. N. et al. (2001). "The Parkes multi-beam pulsar survey–I. Observing and data analysis systems, discovery and timing of 100 pulsars". In: *Monthly Notices of the Royal Astronomical Society* 328.1, pp. 17–35.

Manchester, R. N. et al. (Apr. 2005). "The Australia Telescope National Facility Pulsar Catalogue". In: *AJ* 129.4, pp. 1993–2006. DOI: `10.1086/428488`. arXiv: `astro-ph/0412641 [astro-ph]`.

Mandel, Ilya and Bernhard Müller (Oct. 2020). "Simple recipes for compact remnant masses and natal kicks". In: *MNRAS* 499.3, pp. 3214–3221. DOI: `10.1093/mnras/staa3043`. arXiv: `2006.08360 [astro-ph.HE]`.

Marchetti, T., E. M. Rossi, and A. G. A. Brown (Nov. 2019). "Gaia DR2 in 6D: searching for the fastest stars in the Galaxy". In: *MNRAS* 490.1, pp. 157–171. DOI: 10.1093/mnras/sty2592. arXiv: 1804.10607 [astro-ph.GA].

Margalit, Ben and Brian D. Metzger (Dec. 2017). "Constraining the Maximum Mass of Neutron Stars from Multi-messenger Observations of GW170817". In: *ApJ* 850.2, L19, p. L19. DOI: 10.3847/2041-8213/aa991c. arXiv: 1710.05938 [astro-ph.HE].

Margalit, Ben et al. (Dec. 2018). "Unveiling the engines of fast radio bursts, superluminous supernovae, and gamma-ray bursts". In: *MNRAS* 481.2, pp. 2407–2426. DOI: 10.1093/mnras/sty2417. arXiv: 1806.05690 [astro-ph.HE].

Marsh, T. R. et al. (Sept. 2016). "A radio-pulsing white dwarf binary star". In: *Nature* 537.7620, pp. 374–377. DOI: 10.1038/nature18620. arXiv: 1607.08265 [astro-ph.SR].

Martinez, J. G. et al. (Oct. 2015). "Pulsar J0453+1559: A Double Neutron Star System with a Large Mass Asymmetry". In: *ApJ* 812.2, 143, p. 143. DOI: 10.1088/0004-637X/812/2/143. arXiv: 1509.08805 [astro-ph.HE].

Matthews, Allison M. et al. (Feb. 2016). "The NANOGrav Nine-year Data Set: Astrometric Measurements of 37 Millisecond Pulsars". In: *ApJ* 818.1, 92, p. 92. DOI: 10.3847/0004-637X/818/1/92. arXiv: 1509.08982 [astro-ph.GA].

McLaughlin, M. A. et al. (Feb. 2006). "Transient radio bursts from rotating neutron stars". In: *Nature* 439.7078, pp. 817–820. DOI: 10.1038/nature04440. arXiv: astro-ph/0511587 [astro-ph].

Melrose, D. B. (Dec. 2017). "Coherent emission mechanisms in astrophysical plasmas". In: *Reviews of Modern Plasma Physics* 1.1, 5, p. 5. DOI: 10.1007/s41614-017-0007-0. arXiv: 1707.02009 [physics.plasm-ph].

Menou, Kristen, Rosalba Perna, and Lars Hernquist (Oct. 2001). "Stability and Evolution of Supernova Fallback Disks". In: *ApJ* 559.2, pp. 1032–1046. DOI: 10.1086/322418. arXiv: astro-ph/0102478 [astro-ph].

Mereghetti, Sandro (Dec. 2013). "Pulsars and Magnetars". In: *Brazilian Journal of Physics* 43.5-6, pp. 356–368. DOI: 10.1007/s13538-013-0137-y. arXiv: 1304.4825 [astro-ph.HE].

Metzger, B. D., E. Quataert, and T. A. Thompson (Apr. 2008). "Short-duration gamma-ray bursts with extended emission from protomagnetar spin-down". In: *MNRAS* 385.3, pp. 1455–1460. DOI: 10.1111/j.1365-2966.2008.12923.x. arXiv: 0712.1233 [astro-ph].

Metzger, Brian D., Paz Beniamini, and Dimitrios Giannios (Apr. 2018). "Effects of Fallback Accretion on Protomagnetar Outflows in Gamma-Ray Bursts and Superluminous Supernovae". In: *ApJ* 857.2, 95, p. 95. DOI: 10.3847/1538-4357/aab70c. arXiv: 1802.07750 [astro-ph.HE].

Metzger, Brian D., Edo Berger, and Ben Margalit (May 2017). "Millisecond Magnetar Birth Connects FRB 121102 to Superluminous Supernovae and Long-duration Gamma-Ray Bursts". In: *ApJ* 841.1, 14, p. 14. DOI: 10.3847/1538-4357/aa633d. arXiv: 1701.02370 [astro-ph.HE].

Metzger, Brian D. and Geoffrey C. Bower (Jan. 2014). "Constraints on long-lived remnants of neutron star binary mergers from late-time radio observations of short duration gamma-ray bursts". In: *MNRAS* 437.2, pp. 1821–1827. DOI: 10.1093/mnras/stt2010. arXiv: 1310.4506 [astro-ph.HE].

Metzger, Brian D. et al. (Jan. 2014). "Ionization break-out from millisecond pulsar wind nebulae: an X-ray probe of the origin of superluminous supernovae". In: *MNRAS* 437.1, pp. 703–720. DOI: 10.1093/mnras/stt1922. arXiv: 1307.8115 [astro-ph.HE].

Michel, F. C. and Jr. Goldwire H. C. (Jan. 1970). "Alignment of Oblique Rotators". In: *Astrophys. Lett.* 5, p. 21.

Michel, F. Curtis (June 1988). "Neutron star disk formation from supernova fall-back and possible observational consequences". In: *Nature* 333.6174, pp. 644–645. DOI: 10.1038/333644a0.

Miller, Benjamin et al. (Oct. 2021). "Truncated Marginal Neural Ratio Estimation". In: *Advances in Neural Information Processing Systems* 34, p. 129. DOI: 10.48550/arXiv.2107.01214. arXiv: 2107.01214 [stat.ML].

Miller, M. C. et al. (Dec. 2019). "PSR J0030+0451 Mass and Radius from NICER Data and Implications for the Properties of Neutron Star Matter". In: *ApJ* 887.1, L24, p. L24. DOI: 10.3847/2041-8213/ab50c5. arXiv: 1912.05705 [astro-ph.HE].

Mineshige, S., K. Nomoto, and T. Shigeyama (Jan. 1993). "Viscous-thermal evolution of free accretion disks around new born neutron stars." In: *A&A* 267, pp. 95–100.

Mineshige, Shin et al. (Nov. 1997). "Black Hole Disk Accretion in Supernovae". In: *ApJ* 489.1, pp. 227–233. DOI: 10.1086/304776.

Minsky, Marvin and Seymour Papert (1969). "An introduction to computational geometry". In: *Cambridge tiass., HIT* 479.480, p. 104.

Mishra-Sharma, Siddharth and Kyle Cranmer (Mar. 2022). "Neural simulation-based inference approach for characterizing the Galactic Center γ-ray excess". In: *Phys. Rev. D* 105.6, 063017, p. 063017. DOI: 10.1103/PhysRevD.105.063017. arXiv: 2110.06931 [astro-ph.HE].

Miyamoto, M. and R. Nagai (Jan. 1975). "Three-dimensional models for the distribution of mass in galaxies." In: *PASJ* 27, pp. 533–543.

Morawski, F. and M. Bejger (Oct. 2020). "Neural network reconstruction of the dense matter equation of state derived from the parameters of neutron stars". In: *A&A* 642, A78, A78. DOI: 10.1051/0004-6361/202038130. arXiv: 2006.07194 [astro-ph.HE].

Morris, D. et al. (June 1997). "Pulsar detection at 87GHz." In: *A&A* 322, pp. L17–L20.

Motch, C. et al. (Apr. 2009). "Proper motions of thermally emitting isolated neutron stars measured with Chandra". In: *A&A* 497.2, pp. 423–435. DOI: 10.1051/0004-6361/200811052. arXiv: 0901.1006 [astro-ph.HE].

Mushtukov, Alexander A. et al. (Dec. 2015). "On the maximum accretion luminosity of magnetized neutron stars: connecting X-ray pulsars and ultraluminous X-ray sources". In: *MNRAS* 454.3, pp. 2539–2548. DOI: 10.1093/mnras/stv2087. arXiv: 1506.03600 [astro-ph.HE].

Nagakura, Hiroki, Kohsuke Sumiyoshi, and Shoichi Yamada (Aug. 2019). "Possible Early Linear Acceleration of Proto-neutron Stars via Asymmetric Neutrino Emission in Core-collapse Supernovae". In: *ApJ* 880.2, L28, p. L28. DOI: 10.3847/2041-8213/ab30ca. arXiv: 1907.04863 [astro-ph.HE].

Nandi, Rana and Stefan Schramm (Jan. 2018). "Transport Properties of the Nuclear Pasta Phase with Quantum Molecular Dynamics". In: *ApJ* 852.2, 135, p. 135. DOI: 10.3847/1538-4357/aa9f12. arXiv: 1709.09793 [astro-ph.HE].

Narayan, Ramesh and J. P. Ostriker (Mar. 1990). "Pulsar Populations and Their Evolution". In: *ApJ* 352, p. 222. DOI: 10.1086/168529.

Navarro, Julio F., Carlos S. Frenk, and Simon D. M. White (May 1996). "The Structure of Cold Dark Matter Halos". In: *ApJ* 462, p. 563. DOI: 10.1086/177173. arXiv: astro-ph/9508025 [astro-ph].

Nyquist, H. (July 1928). "Thermal Agitation of Electric Charge in Conductors". In: *Physical Review* 32.1, pp. 110–113. DOI: 10.1103/PhysRev.32.110.

Ofengeim, D. D. and M. E. Gusakov (Aug. 2018). "Fast magnetic field evolution in neutron stars: The key role of magnetically induced fluid motions in the core". In: *Phys. Rev. D* 98.4, 043007, p. 043007. DOI: 10.1103/PhysRevD.98.043007. arXiv: 1805.03956 [astro-ph.HE].

Olausen, S. A. and V. M. Kaspi (May 2014). "The McGill Magnetar Catalog". In: *ApJS* 212.1, 6, p. 6. DOI: 10.1088/0067-0049/212/1/6. arXiv: 1309.4167 [astro-ph.HE].

Oppenheimer, J. R. and G. M. Volkoff (Feb. 1939). "On Massive Neutron Cores". In: *Physical Review* 55.4, pp. 374–381. DOI: 10.1103/PhysRev.55.374.

Osłowski, S. et al. (May 2011). "Population synthesis of double neutron stars". In: *MNRAS* 413.1, pp. 461–479. DOI: 10.1111/j.1365-2966.2010.18147.x. arXiv: 0903.3538 [astro-ph.GA].

Ott, Christian D. et al. (May 2006). "The Spin Periods and Rotational Profiles of Neutron Stars at Birth". In: *ApJS* 164.1, pp. 130–155. DOI: 10.1086/500832. arXiv: astro-ph/0508462 [astro-ph].

Özel, Feryal and Paulo Freire (Sept. 2016). "Masses, Radii, and the Equation of State of Neutron Stars". In: *ARA&A* 54, pp. 401–440. DOI: 10.1146/annurev-astro-081915-023322. arXiv: 1603.02698 [astro-ph.HE].

Pacini, F. (July 1968). "Rotating Neutron Stars, Pulsars and Supernova Remnants". In: *Nature* 219.5150, pp. 145–146. DOI: 10.1038/219145a0.

Papamakarios, George and Iain Murray (May 2016). "Fast ϵ-free Inference of Simulation Models with Bayesian Conditional Density Estimation". In: *arXiv e-prints*, arXiv:1605.06376, arXiv:1605.06376. DOI: 10.48550/arXiv.1605.06376. arXiv: 1605.06376 [stat.ML].

Papamakarios, George, Theo Pavlakou, and Iain Murray (May 2017). "Masked Autoregressive Flow for Density Estimation". In: *arXiv e-prints*, arXiv:1705.07057, arXiv:1705.07057. DOI: 10.48550/arXiv.1705.07057. arXiv: 1705.07057 [stat.ML].

Papamakarios, George, David C. Sterratt, and Iain Murray (May 2018). "Sequential Neural Likelihood: Fast Likelihood-free Inference with Autoregressive Flows". In: *arXiv e-prints*, arXiv:1805.07226, arXiv:1805.07226. DOI: 10.48550/arXiv.1805.07226. arXiv: 1805.07226 [stat.ML].

Parfrey, Kyle, Anatoly Spitkovsky, and Andrei M. Beloborodov (May 2016). "Torque Enhancement, Spin Equilibrium, and Jet Power from Disk-Induced Opening of Pulsar Magnetic Fields". In: *ApJ* 822.1, 33, p. 33. DOI: 10.3847/0004-637X/822/1/33. arXiv: 1507.08627 [astro-ph.HE].

Parthasarathy, A. et al. (Nov. 2019). "Timing of young radio pulsars - I. Timing noise, periodic modulation, and proper motion". In: *MNRAS* 489.3, pp. 3810–3826. DOI: 10.1093/mnras/stz2383. arXiv: 1908.11709 [astro-ph.HE].

Passamonti, Andrea et al. (Aug. 2017). "On the magnetic field evolution time-scale in superconducting neutron star cores". In: *MNRAS* 469.4, pp. 4979–4984. DOI: 10.1093/mnras/stx1192. arXiv: 1704.02016 [astro-ph.HE].

Paszke, Adam et al. (Dec. 2019). "PyTorch: An Imperative Style, High-Performance Deep Learning Library". In: *preprint*, arXiv:1912.01703, arXiv:1912.01703. arXiv: 1912.01703 [cs.LG].

Pavan, L. et al. (Feb. 2014). "The long helical jet of the Lighthouse nebula, IGR J11014-6103". In: *A&A* 562, A122, A122. DOI: 10.1051/0004-6361/201322588. arXiv: 1309.6792 [astro-ph.HE].

Pelisoli, Ingrid et al. (June 2023). "A 5.3-min-period pulsing white dwarf in a binary detected from radio to X-rays". In: *Nature Astronomy*. DOI: 10.1038/s41550-023-01995-x. arXiv: 2306.09272 [astro-ph.SR].

Perera, B. B. P. et al. (Dec. 2019). "The International Pulsar Timing Array: second data release". In: *MNRAS* 490.4, pp. 4666–4687. DOI: 10.1093/mnras/stz2857. arXiv: 1909.04534 [astro-ph.HE].

Perna, Rosalba et al. (Feb. 2014). "The Fate of Fallback Matter around Newly Born Compact Objects". In: *ApJ* 781.2, 119, p. 119. DOI: 10.1088/0004-637X/781/2/119. arXiv: 1312.4981 [astro-ph.HE].

Pétri, Jérôme (Oct. 2016). "Theory of pulsar magnetosphere and wind". In: *Journal of Plasma Physics* 82.5, 635820502, p. 635820502. DOI: 10.1017/S0022377816000763. arXiv: 1608.04895 [astro-ph.HE].

— (Jan. 2020). "Electrodynamics and Radiation from Rotating Neutron Star Magnetospheres". In: *Universe* 6.1, p. 15. DOI: 10.3390/universe6010015. arXiv: 2001.03422 [astro-ph.HE].

Petroff, E., J. W. T. Hessels, and D. R. Lorimer (Dec. 2022). "Fast radio bursts at the dawn of the 2020s". In: *A&ARv* 30.1, 2, p. 2. DOI: 10.1007/s00159-022-00139-w. arXiv: 2107.10113 [astro-ph.HE].

Philippov, Alexander, Alexander Tchekhovskoy, and Jason G. Li (July 2014). "Time evolution of pulsar obliquity angle from 3D simulations of magnetospheres". In: *MNRAS* 441.3, pp. 1879–1887. DOI: 10.1093/mnras/stu591. arXiv: 1311.1513 [astro-ph.HE].

Philippov, Alexander, Andrey Timokhin, and Anatoly Spitkovsky (June 2020). "Origin of Pulsar Radio Emission". In: *Phys. Rev. Lett.* 124.24, 245101, p. 245101. DOI: 10.1103/PhysRevLett.124.245101. arXiv: 2001.02236 [astro-ph.HE].

Phillips, A. C. (1994). *The physics of stars*.

Phinney, E. S. and R. D. Blandford (Jan. 1981). "Analysis of the pulsar P-P-prime distribution". In: *MNRAS* 194, pp. 137–148. DOI: 10.1093/mnras/194.1.137.

Pichardo, Bárbara et al. (Mar. 2012). "The Sun was Not Born in M67". In: *AJ* 143.3, 73, p. 73. DOI: 10.1088/0004-6256/143/3/73. arXiv: 1201.0987 [astro-ph.GA].

Piro, Anthony L. and Christian D. Ott (Aug. 2011). "Supernova Fallback onto Magnetars and Propeller-powered Supernovae". In: *ApJ* 736.2, 108, p. 108. DOI: 10.1088/0004-637X/736/2/108. arXiv: 1104.0252 [astro-ph.HE].

Platts, E. et al. (Aug. 2019). "A living theory catalogue for fast radio bursts". In: *Phys. Rep.* 821, pp. 1–27. DOI: 10.1016/j.physrep.2019.06.003. arXiv: 1810.05836 [astro-ph.HE].

Podsiadlowski, Ph. et al. (Sept. 2004). "The Effects of Binary Evolution on the Dynamics of Core Collapse and Neutron Star Kicks". In: *ApJ* 612.2, pp. 1044–1051. DOI: 10.1086/421713. arXiv: astro-ph/0309588 [astro-ph].

Polido, P., F. Jablonski, and J. R. D. Lépine (Nov. 2013). "A Galaxy Model from Two Micron All Sky Survey Star Counts in the Whole Sky, Including the Plane". In: *ApJ* 778.1, 32, p. 32. DOI: 10.1088/0004-637X/778/1/32. arXiv: 1308.6238 [astro-ph.GA].

Pons, J. A. and U. Geppert (July 2007). "Magnetic field dissipation in neutron star crusts: from magnetars to isolated neutron stars". In: *A&A* 470.1, pp. 303–315. DOI: 10.1051/0004-6361:20077456. arXiv: astro-ph/0703267 [astro-ph].

Pons, J. A., J. A. Miralles, and U. Geppert (Mar. 2009). "Magneto-thermal evolution of neutron stars". In: *A&A* 496.1, pp. 207–216. DOI: 10.1051/0004-6361:200811229. arXiv: 0812.3018 [astro-ph].

Pons, José A. and Daniele Viganò (Dec. 2019). "Magnetic, thermal and rotational evolution of isolated neutron stars". In: *Living Reviews in Computational Astrophysics* 5.1, 3, p. 3. DOI: 10.1007/s41115-019-0006-7. arXiv: 1911.03095 [astro-ph.HE].

Pons, José A., Daniele Viganò, and Nanda Rea (July 2013). "A highly resistive layer within the crust of X-ray pulsars limits their spin periods". In: *Nature Physics* 9.7, pp. 431–434. DOI: 10.1038/nphys2640. arXiv: 1304.6546 [astro-ph.SR].

Popov, S. B. and R. Turolla (Oct. 2012). "Initial spin periods of neutron stars in supernova remnants". In: *Ap&SS* 341.2, pp. 457–464. DOI: 10.1007/s10509-012-1100-z. arXiv: 1204.0632 [astro-ph.HE].

Popov, S. B. et al. (Feb. 2010). "Population synthesis studies of isolated neutron stars with magnetic field decay". In: *MNRAS* 401.4, pp. 2675–2686. DOI: 10.1111/j.1365-2966.2009.15850.x. arXiv: 0910.2190 [astro-ph.HE].

Posselt, B. et al. (Sept. 2018). "Discovery of Extended Infrared Emission around the Neutron Star RXJ0806.4-4123". In: *ApJ* 865.1, 1, p. 1. DOI: 10.3847/1538-4357/aad6df. arXiv: 1809.08107 [astro-ph.HE].

Posselt, B. et al. (Apr. 2023). "The Thousand Pulsar Array program on MeerKAT - IX. The time-averaged properties of the observed pulsar population". In: *MNRAS* 520.3, pp. 4582–4600. DOI: 10.1093/mnras/stac3383. arXiv: 2211.11849 [astro-ph.HE].

Potekhin, Alexander Y., José A. Pons, and Dany Page (Oct. 2015). "Neutron Stars—Cooling and Transport". In: *Space Sci. Rev.* 191.1-4, pp. 239–291. DOI: 10.1007/s11214-015-0180-9. arXiv: 1507.06186 [astro-ph.HE].

Poutanen, Juri et al. (May 2007). "Supercritically accreting stellar mass black holes as ultraluminous X-ray sources". In: *MNRAS* 377.3, pp. 1187–1194. DOI: 10.1111/j.1365-2966.2007.11668.x. arXiv: astro-ph/0609274 [astro-ph].

Prasanna, Tejas et al. (Dec. 2023). "The early evolution of magnetar rotation - II. Rapidly rotating magnetars: implications for gamma-ray bursts and superluminous supernovae". In: *MNRAS* 526.2, pp. 3141–3155. DOI: 10.1093/mnras/stad2948. arXiv: 2305.16412 [astro-ph.HE].

Prialnik, Dina (2000). *An Introduction to the Theory of Stellar Structure and Evolution*.

Quinlan, J. R. (1986). "Induction of Decision Trees". In: *Machine Learning* 1, pp. 81–106.

Raaijmakers, G. et al. (Sept. 2021). "Constraints on the Dense Matter Equation of State and Neutron Star Properties from NICER's Mass-Radius Estimate of PSR J0740+6620 and Multimessenger Observations". In: *ApJ* 918.2, L29, p. L29. DOI: 10.3847/2041-8213/ac089a. arXiv: 2105.06981 [astro-ph.HE].

Rankin, Joanna M. (Mar. 1993). "Toward an Empirical Theory of Pulsar Emission. VI. The Geometry of the Conal Emission Region". In: *ApJ* 405, p. 285. DOI: 10.1086/172361.

Rawat, Waseem and Zenghui Wang (June 2017). "Deep Convolutional Neural Networks for Image Classification: A Comprehensive Review". In: *Neural Computation* 29, pp. 1–98. DOI: 10.1162/NECO_a_00990.

Re Fiorentin, P. et al. (June 2007). "Estimation of stellar atmospheric parameters from SDSS/SEGUE spectra". In: *A&A* 467.3, pp. 1373–1387. DOI: 10.1051/0004-6361:20077334. arXiv: astro-ph/0703309 [astro-ph].

Rea, N. et al. (Sept. 2016). "Magnetar-like Activity from the Central Compact Object in the SNR RCW103". In: *ApJ* 828.1, L13, p. L13. DOI: 10.3847/2041-8205/828/1/L13. arXiv: 1607.04107 [astro-ph.HE].

Rea, N. et al. (Nov. 2022). "Constraining the Nature of the 18 min Periodic Radio Transient GLEAM-X J162759.5-523504.3 via Multiwavelength Observations and Magneto-thermal Simulations". In: *ApJ* 940.1, 72, p. 72. DOI: 10.3847/1538-4357/ac97ea. arXiv: 2210.01903 [astro-ph.HE].

Rea, Nanda et al. (July 2023). "A long-period radio transient active for three decades: population study in the neutron star and white dwarf rotating dipole scenarios". In: *arXiv e-prints*, arXiv:2307.10351, arXiv:2307.10351. DOI: 10.48550/arXiv.2307.10351. arXiv: 2307.10351 [astro-ph.HE].

Reed, B. Cameron (Oct. 2005). "New Estimates of the Solar-Neighborhood Massive Star Birthrate and the Galactic Supernova Rate". In: *AJ* 130.4, pp. 1652–1657. DOI: 10.1086/444474. arXiv: astro-ph/0506708 [astro-ph].

Reid, M. J. et al. (Nov. 2019). "Trigonometric Parallaxes of High-mass Star-forming Regions: Our View of the Milky Way". In: *ApJ* 885.2, 131, p. 131. DOI: 10.3847/1538-4357/ab4a11. arXiv: 1910.03357 [astro-ph.GA].

Remazeilles, M. et al. (Aug. 2015). "An improved source-subtracted and destriped 408-MHz all-sky map". In: *MNRAS* 451.4, pp. 4311–4327. DOI: 10.1093/mnras/stv1274. arXiv: 1411.3628 [astro-ph.IM].

Richards, D. W. and J. M. Comella (May 1969). "The Period of Pulsar NP 0532". In: *Nature* 222.5193, pp. 551–552. DOI: 10.1038/222551a0.

Riley, T. E. et al. (Dec. 2019). "A NICER View of PSR J0030+0451: Millisecond Pulsar Parameter Estimation". In: *ApJ* 887.1, L21, p. L21. DOI: 10.3847/2041-8213/ab481c. arXiv: 1912.05702 [astro-ph.HE].

Riley, Thomas E. et al. (Sept. 2021). "A NICER View of the Massive Pulsar PSR J0740+6620 Informed by Radio Timing and XMM-Newton Spectroscopy". In: *ApJ* 918.2, L27, p. L27. DOI: 10.3847/2041-8213/ac0a81. arXiv: 2105.06980 [astro-ph.HE].

Romani, Roger W. et al. (Aug. 2022). "PSR J0952-0607: The Fastest and Heaviest Known Galactic Neutron Star". In: *ApJ* 934.2, L17, p. L17. DOI: 10.3847/2041-8213/ac8007. arXiv: 2207.05124 [astro-ph.HE].

Romanova, M. M. et al. (Dec. 2005). "Propeller-driven Outflows and Disk Oscillations". In: *ApJ* 635.2, pp. L165–L168. DOI: 10.1086/499560. arXiv: astro-ph/0511600 [astro-ph].

Ronchi, M. et al. (Aug. 2021). "Analyzing the Galactic Pulsar Distribution with Machine Learning". In: *ApJ* 916.2, 100, p. 100. DOI: 10.3847/1538-4357/ac05bd. arXiv: 2101.06145 [astro-ph.HE].

Ronchi, M. et al. (Aug. 2022). "Long-period Pulsars as Possible Outcomes of Supernova Fallback Accretion". In: *ApJ* 934.2, 184, p. 184. DOI: 10.3847/1538-4357/ac7cec. arXiv: 2201.11704 [astro-ph.HE].

Rosenblatt, F. (1958). "The perceptron: A probabilistic model for information storage and organization in the brain." In: *Psychological Review* 65.6, pp. 386–408. ISSN: 0033-295X. DOI: 10.1037/h0042519. URL: http://dx.doi.org/10.1037/h0042519.

Rozwadowska, Karolina, Francesco Vissani, and Enrico Cappellaro (Feb. 2021). "On the rate of core collapse supernovae in the milky way". In: *New Astron.* 83, 101498, p. 101498. DOI: 10.1016/j.newast.2020.101498. arXiv: 2009.03438 [astro-ph.HE].

Rubin, Donal B. (Dec. 1984). "Bayesianly Justifiable and Relevant Frequency Calculations for the Applied Statistician". In: *Ann. Statist.* 12.4, pp. 1151–1172. DOI: 10.1214/aos/1176346785. URL: https://doi.org/10.1214/aos/1176346785.

Rudak, B. and H. Ritter (Apr. 1994). "The Line of Death - the Line of Birth". In: *MNRAS* 267, p. 513. DOI: 10.1093/mnras/267.3.513.

Ruder, Sebastian (2017). *An overview of gradient descent optimization algorithms*. arXiv: 1609.04747 [cs.LG].

Ruderman, M. A. and P. G. Sutherland (Feb. 1975). "Theory of pulsars: polar gaps, sparks, and coherent microwave radiation." In: *ApJ* 196, pp. 51–72. DOI: 10.1086/153393.

Rybicki, George B. and Alan P. Lightman (1986). *Radiative Processes in Astrophysics*.

Sánchez Almeida, J. and C. Allende Prieto (Jan. 2013). "Automated Unsupervised Classification of the Sloan Digital Sky Survey Stellar Spectra using k-means Clustering". In: *ApJ* 763.1, 50, p. 50. DOI: 10.1088/0004-637X/763/1/50. arXiv: 1211.5321 [astro-ph.SR].

Sartore, N. et al. (Feb. 2010). "Galactic neutron stars. I. Space and velocity distributions in the disk and in the halo". In: *A&A* 510, A23, A23. DOI: 10.1051/0004-6361/200912222. arXiv: 0908.3182 [astro-ph.GA].

Schwab, J., Ph. Podsiadlowski, and S. Rappaport (Aug. 2010). "Further Evidence for the Bimodal Distribution of Neutron-star Masses". In: *ApJ* 719.1, pp. 722–727. DOI: 10.1088/0004-637X/719/1/722. arXiv: 1006.4584 [astro-ph.HE].

Shabaltas, Natalia and Dong Lai (Apr. 2012). "The Hidden Magnetic Field of the Young Neutron Star in Kesteven 79". In: *ApJ* 748.2, 148, p. 148. DOI: 10.1088/0004-637X/748/2/148. arXiv: 1110.3129 [astro-ph.HE].

Shapiro, Stuart L. and Saul A. Teukolsky (1983). *Black holes, white dwarfs and neutron stars. The physics of compact objects*. DOI: 10.1002/9783527617661.

Sharma, Sanjib (Aug. 2017). "Markov Chain Monte Carlo Methods for Bayesian Data Analysis in Astronomy". In: *ARA&A* 55.1, pp. 213–259. DOI: 10.1146/annurev-astro-082214-122339. arXiv: 1706.01629 [astro-ph.IM].

Shklovskii, I. S. (Feb. 1970). "Possible Causes of the Secular Increase in Pulsar Periods." In: *Soviet Ast.* 13, p. 562.

Skliris, Vasileios, Michael R. K. Norman, and Patrick J. Sutton (Sept. 2020). "Real-Time Detection of Unmodeled Gravitational-Wave Transients UsingConvolutional Neural Networks". In: *preprint*, arXiv:2009.14611, arXiv:2009.14611. arXiv: 2009.14611 [astro-ph.IM].

Skowron, Dorota M. et al. (Aug. 2019). "A three-dimensional map of the Milky Way using classical Cepheid variable stars". In: *Science* 365.6452, pp. 478–482. DOI: 10.1126/science.aau3181. arXiv: 1806.10653 [astro-ph.GA].

Skrzypczak, Anna et al. (Feb. 2018). "Meterwavelength Single-pulse Polarimetric Emission Survey. IV. The Period Dependence of Component Widths of Pulsars". In: *ApJ* 854.2, 162, p. 162. DOI: 10.3847/1538-4357/aaa758. arXiv: 1801.02720 [astro-ph.HE].

Smith, David A. et al. (July 2023). "The Third Fermi Large Area Telescope Catalog of Gamma-ray Pulsars". In: *arXiv e-prints*, arXiv:2307.11132, arXiv:2307.11132. DOI: 10.48550/arXiv.2307.11132. arXiv: 2307.11132 [astro-ph.HE].

Smits, R. et al. (Jan. 2009). "Pulsar searches and timing with the square kilometre array". In: *A&A* 493.3, pp. 1161–1170. DOI: 10.1051/0004-6361:200810383. arXiv: 0811.0211 [astro-ph].

Smits, R. et al. (Apr. 2011). "Prospects for accurate distance measurements of pulsars with the Square Kilometre Array: Enabling fundamental physics". In: *A&A* 528, A108, A108. DOI: 10.1051/0004-6361/201016141. arXiv: 1101.5971 [astro-ph.IM].

Solanes, Aleix and Joaquim Radua (Apr. 2022). "Advances in Using MRI to Estimate the Risk of Future Outcomes in Mental Health - Are We Getting There?" In: *Frontiers in Psychiatry* 13, fpsyt–13. DOI: 10.3389/fpsyt.2022.826111.

Speagle, Joshua S. (Apr. 2020). "DYNESTY: a dynamic nested sampling package for estimating Bayesian posteriors and evidences". In: *MNRAS* 493.3, pp. 3132–3158. DOI: 10.1093/mnras/staa278. arXiv: 1904.02180 [astro-ph.IM].

Spitkovsky, Anatoly (Sept. 2006). "Time-dependent Force-free Pulsar Magnetospheres: Axisymmetric and Oblique Rotators". In: *ApJ* 648.1, pp. L51–L54. DOI: 10.1086/507518. arXiv: astro-ph/0603147 [astro-ph].

Spitzer Lyman, Jr. (May 1942). "The Dynamics of the Interstellar Medium. III. Galactic Distribution." In: *ApJ* 95, p. 329. DOI: 10.1086/144407.

Springer, M.D. (1979). *The Algebra of Random Variables*. Probability and Statistics Series. Wiley. ISBN: 9780471014065. URL: https://books.google.es/books?id=qUDvAAAAMAAJ.

Staelin, David H. and III Reifenstein Edward C. (Dec. 1968). "Pulsating Radio Sources near the Crab Nebula". In: *Science* 162.3861, pp. 1481–1483. DOI: 10.1126/science.162.3861.1481.

Staubert, R. et al. (Feb. 2019). "Cyclotron lines in highly magnetized neutron stars". In: *A&A* 622, A61, A61. DOI: 10.1051/0004-6361/201834479. arXiv: 1812.03461 [astro-ph.HE].

Steiner, A. W. et al. (May 2018). "Constraining the mass and radius of neutron stars in globular clusters". In: *MNRAS* 476.1, pp. 421–435. DOI: 10.1093/mnras/sty215. arXiv: 1709.05013 [astro-ph.HE].

Stovall, K. et al. (Aug. 2014). "The Green Bank Northern Celestial Cap Pulsar Survey. I. Survey Description, Data Analysis, and Initial Results". In: *ApJ* 791.1, 67, p. 67. DOI: 10.1088/0004-637X/791/1/67. arXiv: 1406.5214 [astro-ph.HE].

Stuart, Alan and J. Keith Ord (1994). *Kendall's advanced theory of statistics. Vol.1: Distribution theory*.

Sturrock, P. A. (Mar. 1971). "A Model of Pulsars". In: *ApJ* 164, p. 529. DOI: 10.1086/150865.

Suvorov, Arthur G. and Andrew Melatos (Mar. 2023). "Evolutionary implications of a magnetar interpretation for GLEAM-X J162759.5-523504.3". In: *MNRAS* 520.1, pp. 1590–1600. DOI: 10.1093/mnras/stad274. arXiv: 2301.08541 [astro-ph.HE].

Szary, Andrzej et al. (Mar. 2014). "Radio Efficiency of Pulsars". In: *ApJ* 784.1, 59, p. 59. DOI: 10.1088/0004-637X/784/1/59. arXiv: 1402.0228 [astro-ph.HE].

Taam, R. E. and E. P. J. van den Heuvel (June 1986). "Magnetic Field Decay and the Origin of Neutron Star Binaries". In: *ApJ* 305, p. 235. DOI: 10.1086/164243.

Talts, Sean et al. (Apr. 2018). "Validating Bayesian Inference Algorithms with Simulation-Based Calibration". In: *arXiv e-prints*, arXiv:1804.06788, arXiv:1804.06788. DOI: 10.48550/arXiv.1804.06788. arXiv: 1804.06788 [stat.ME].

Tamborra, Irene et al. (Sept. 2014). "Self-sustained Asymmetry of Lepton-number Emission: A New Phenomenon during the Supernova Shock-accretion Phase in Three Dimensions". In: *ApJ* 792.2, 96, p. 96. DOI: 10.1088/0004-637X/792/2/96. arXiv: 1402.5418 [astro-ph.SR].

Tan, C. M. et al. (Oct. 2018). "LOFAR Discovery of a 23.5 s Radio Pulsar". In: *ApJ* 866.1, 54, p. 54. DOI: 10.3847/1538-4357/aade88. arXiv: 1809.00965 [astro-ph.HE].

Tauris, T. M., N. Langer, and M. Kramer (Sept. 2012). "Formation of millisecond pulsars with CO white dwarf companions - II. Accretion, spin-up, true ages and comparison to MSPs with He white dwarf companions". In: *MNRAS* 425.3, pp. 1601–1627. DOI: 10.1111/j.1365-2966.2012.21446.x. arXiv: 1206.1862 [astro-ph.SR].

Tauris, T. M. and R. N. Manchester (Aug. 1998). "On the Evolution of Pulsar Beams". In: *MNRAS* 298.3, pp. 625–636. DOI: 10.1046/j.1365-8711.1998.01369.x.

Tauris, T. M. et al. (Dec. 2013). "Ultra-stripped Type Ic Supernovae from Close Binary Evolution". In: *ApJ* 778.2, L23, p. L23. DOI: 10.1088/2041-8205/778/2/L23. arXiv: 1310.6356 [astro-ph.SR].

Tauris, T. M. et al. (Apr. 2015). "Understanding the Neutron Star Population with the SKA". In: *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, p. 39. arXiv: 1501.00005 [astro-ph.HE].

Tauris, Thomas M. and Gerrit J. Savonije (Oct. 1999). "Formation of millisecond pulsars. I. Evolution of low-mass X-ray binaries with P_orb> 2 days". In: *A&A* 350,

pp. 928–944. DOI: `10.48550/arXiv.astro-ph/9909147`. arXiv: `astro-ph/9909147 [astro-ph]`.

Taylor, J. H. and J. M. Cordes (July 1993). "Pulsar Distances and the Galactic Distribution of Free Electrons". In: *ApJ* 411, p. 674. DOI: `10.1086/172870`.

Tejero-Cantero, Alvaro et al. (Aug. 2020). "sbi: A toolkit for simulation-based inference". In: *The Journal of Open Source Software* 5.52, 2505, p. 2505. DOI: `10.21105/joss.02505`.

Tenenbaum, J. B., V. Silva, and J. C. Langford (2000). "A Global Geometric Framework for Nonlinear Dimensionality Reduction". In: *Science* 290.5500, pp. 2319–2323.

Teo, Yong et al. (Oct. 2021). "Benchmarking quantum tomography completeness and fidelity with machine learning". In: *New Journal of Physics* 23. DOI: `10.1088/1367-2630/ac1fcb`.

Thompson, Christopher and Robert C. Duncan (July 1995). "The soft gamma repeaters as very strongly magnetized neutron stars - I. Radiative mechanism for outbursts". In: *MNRAS* 275.2, pp. 255–300. DOI: `10.1093/mnras/275.2.255`.

Thompson, Christopher et al. (Nov. 2000). "Physical Mechanisms for the Variable Spin-down and Light Curve of SGR 1900+14". In: *ApJ* 543.1, pp. 340–350. DOI: `10.1086/317072`. arXiv: `astro-ph/9908086 [astro-ph]`.

Thompson, Todd A., Philip Chang, and Eliot Quataert (Aug. 2004). "Magnetar Spin-Down, Hyperenergetic Supernovae, and Gamma-Ray Bursts". In: *ApJ* 611.1, pp. 380–393. DOI: `10.1086/421969`. arXiv: `astro-ph/0401555 [astro-ph]`.

Tingay, S. J. et al. (Jan. 2013). "The Murchison Widefield Array: The Square Kilometre Array Precursor at Low Radio Frequencies". In: *Publ. Astron. Soc. Australia* 30, e007, e007. DOI: `10.1017/pasa.2012.007`. arXiv: `1206.6945 [astro-ph.IM]`.

Tolman, Richard C. (Feb. 1939). "Static Solutions of Einstein's Field Equations for Spheres of Fluid". In: *Physical Review* 55.4, pp. 364–373. DOI: `10.1103/PhysRev.55.364`.

Tong, H. (Apr. 2022). "Discussions on the nature of GLEAM-X J162759.5-523504.3". In: *arXiv e-prints*, arXiv:2204.01957, arXiv:2204.01957. arXiv: `2204.01957 [astro-ph.HE]`.

Tong, H. et al. (May 2013). "Wind Braking of Magnetars". In: *ApJ* 768.2, 144, p. 144. DOI: `10.1088/0004-637X/768/2/144`. arXiv: `1205.1626 [astro-ph.HE]`.

Tong, H. et al. (Dec. 2016). "Rotational Evolution of Magnetars in the Presence of a Fallback Disk". In: *ApJ* 833.2, 265, p. 265. DOI: `10.3847/1538-4357/833/2/265`. arXiv: `1608.02113 [astro-ph.HE]`.

Torres-Forné, Alejandro et al. (Mar. 2016). "Are pulsars born with a hidden magnetic field?" In: *MNRAS* 456.4, pp. 3813–3826. DOI: `10.1093/mnras/stv2926`. arXiv: `1511.03823 [astro-ph.SR]`.

Tout, C. A. et al. (June 2008). "Binary star origin of high field magnetic white dwarfs". In: *MNRAS* 387.2, pp. 897–901. DOI: `10.1111/j.1365-2966.2008.13291.x`. arXiv: `0805.0115 [astro-ph]`.

Turolla, R., S. Zane, and A. L. Watts (Nov. 2015). "Magnetars: the physics behind observations. A review". In: *Reports on Progress in Physics* 78.11, 116901, p. 116901. DOI: `10.1088/0034-4885/78/11/116901`. arXiv: `1507.02924 [astro-ph.HE]`.

Turolla, Roberto (Jan. 2009). "Isolated Neutron Stars: The Challenge of Simplicity". In: *Astrophysics and Space Science Library*. Ed. by Werner Becker. Vol. 357. Astrophysics and Space Science Library, p. 141. DOI: `10.1007/978-3-540-76965-1_7`.

Ugliano, Marcella et al. (Sept. 2012). "Progenitor-explosion Connection and Remnant Birth Masses for Neutrino-driven Supernovae of Iron-core Progenitors". In: *ApJ* 757.1, 69, p. 69. DOI: `10.1088/0004-637X/757/1/69`. arXiv: `1205.3657 [astro-ph.SR]`.

Vallée, Jacques P. (2017). "A guided map to the spiral arms in the galactic disk of the Milky Way". In: *Astron. Rev.* 13.3-4, pp. 113–146. DOI: 10.1080/21672857.2017.1379459. eprint: https://doi.org/10.1080/21672857.2017.1379459. URL: https://doi.org/10.1080/21672857.2017.1379459.

van Haarlem, M. P. et al. (Aug. 2013). "LOFAR: The LOw-Frequency ARray". In: *A&A* 556, A2, A2. DOI: 10.1051/0004-6361/201220873. arXiv: 1305.3550 [astro-ph.IM].

van Kerkwijk, M. H. and D. L. Kaplan (Apr. 2007). "Isolated neutron stars: magnetic fields, distances,and spectra". In: *Ap&SS* 308.1-4, pp. 191–201. DOI: 10.1007/s10509-007-9343-9. arXiv: astro-ph/0607320 [astro-ph].

Vasist, Malavika et al. (Apr. 2023). "Neural posterior estimation for exoplanetary atmospheric retrieval". In: *A&A* 672, A147, A147. DOI: 10.1051/0004-6361/202245263. arXiv: 2301.06575 [astro-ph.EP].

Verbunt, Frank, Andrei Igoshev, and Eric Cator (Dec. 2017). "The observed velocity distribution of young pulsars". In: *A&A* 608, A57, A57. DOI: 10.1051/0004-6361/201731518. arXiv: 1708.08281 [astro-ph.HE].

Viganò, D. and J. A. Pons (Oct. 2012). "Central compact objects and the hidden magnetic field scenario". In: *MNRAS* 425.4, pp. 2487–2492. DOI: 10.1111/j.1365-2966.2012.21679.x. arXiv: 1206.2014 [astro-ph.SR].

Viganò, D., J. A. Pons, and J. A. Miralles (Oct. 2012). "A new code for the Hall-driven magnetic evolution of neutron stars". In: *Computer Physics Communications* 183.10, pp. 2042–2053. DOI: 10.1016/j.cpc.2012.04.029. arXiv: 1204.4707 [astro-ph.SR].

Viganò, D. et al. (Sept. 2013). "Unifying the observational diversity of isolated neutron stars via magneto-thermal evolution models". In: *MNRAS* 434.1, pp. 123–141. DOI: 10.1093/mnras/stt1008. arXiv: 1306.2156 [astro-ph.SR].

Viganò, Daniele and Diego F. Torres (June 2015). "Modelling of the γ-ray pulsed spectra of Geminga, Crab, and Vela with synchro-curvature radiation". In: *MNRAS* 449.4, pp. 3755–3765. DOI: 10.1093/mnras/stv579. arXiv: 1503.04060 [astro-ph.HE].

Viganò, Daniele, Diego F. Torres, and Jonatan Martín (Nov. 2015). "A systematic synchro-curvature modelling of pulsar γ-ray spectra unveils hidden trends". In: *MNRAS* 453.3, pp. 2599–2621. DOI: 10.1093/mnras/stv1582. arXiv: 1507.04021 [astro-ph.HE].

Viganò, Daniele et al. (Aug. 2021). "Magneto-thermal evolution of neutron stars with coupled Ohmic, Hall and ambipolar effects via accurate finite-volume simulations". In: *Computer Physics Communications* 265, 108001, p. 108001. DOI: 10.1016/j.cpc.2021.108001. arXiv: 2104.08001 [astro-ph.HE].

Vivekanand, M. and R. Narayan (Sept. 1981). "A new look at pulsar statistics - Birthrate and evidence for injection". In: *Journal of Astrophysics and Astronomy* 2, pp. 315–337. DOI: 10.1007/BF02714556.

Vivekanand, M., R. Narayan, and V. Radhakrishnan (Sept. 1982). "On selection effects in pulsar searches". In: *Journal of Astrophysics and Astronomy* 3, pp. 237–247. DOI: 10.1007/BF02714864.

Wainscoat, Richard J. et al. (Nov. 1992). "A Model of the 8–25 Micron Point Source Infrared Sky". In: *ApJS* 83, p. 111. DOI: 10.1086/191733.

Walter, F. M. et al. (Nov. 2010). "Revisiting the Parallax of the Isolated Neutron Star RX J185635-3754 Using HST/ACS Imaging". In: *ApJ* 724.1, pp. 669–677. DOI: 10.1088/0004-637X/724/1/669. arXiv: 1008.1709 [astro-ph.SR].

Wang, J. B. et al. (July 2017). "Comparison of pulsar positions from timing and very long baseline astrometry". In: *MNRAS* 469.1, pp. 425–434. DOI: 10.1093/mnras/stx837. arXiv: 1704.01011 [astro-ph.HE].

Wang, Zhongxiang, Deepto Chakrabarty, and David L. Kaplan (Apr. 2006). "A debris disk around an isolated young neutron star". In: *Nature* 440.7085, pp. 772–775. DOI: 10.1038/nature04669. arXiv: astro-ph/0604076 [astro-ph].

Wayth, Randall B. et al. (Nov. 2018). "The Phase II Murchison Widefield Array: Design overview". In: *Publ. Astron. Soc. Australia* 35, e033, e033. DOI: 10.1017/pasa.2018.37. arXiv: 1809.06466 [astro-ph.IM].

Wei, Wei and E. A. Huerta (Oct. 2020). "Deep learning for gravitational wave forecasting of neutron star mergers". In: *preprint*, arXiv:2010.09751, arXiv:2010.09751. arXiv: 2010.09751 [gr-qc].

Wheeler, J. Craig et al. (July 2000). "Asymmetric Supernovae, Pulsars, Magnetars, and Gamma-Ray Bursts". In: *ApJ* 537.2, pp. 810–823. DOI: 10.1086/309055. arXiv: astro-ph/9909293 [astro-ph].

White, Christopher J. et al. (Feb. 2022). "On the Origin of Pulsar and Magnetar Magnetic Fields". In: *ApJ* 926.2, 111, p. 111. DOI: 10.3847/1538-4357/ac4507. arXiv: 2111.01814 [astro-ph.HE].

Wickramasinghe, D. T. and Lilia Ferrario (July 2000). "Magnetism in Isolated and Binary White Dwarfs". In: *PASP* 112.773, pp. 873–924. DOI: 10.1086/316593.

Wood, Toby S. and Vanessa Graber (Apr. 2022). "Superconducting Phases in Neutron Star Cores". In: *Universe* 8.4, p. 228. DOI: 10.3390/universe8040228.

Woods, P. M. and C. Thompson (2006). "Soft gamma repeaters and anomalous X-ray pulsars: magnetar candidates". In: *Compact stellar X-ray sources*. Vol. 39, pp. 547–586. DOI: 10.48550/arXiv.astro-ph/0406133.

Woosley, S. E. (Aug. 2010). "Bright Supernovae from Magnetar Birth". In: *ApJ* 719.2, pp. L204–L207. DOI: 10.1088/2041-8205/719/2/L204. arXiv: 0911.0698 [astro-ph.HE].

Wu, Qing-Dong et al. (Nov. 2020). "Luminosity of a radio pulsar and its new emission death line". In: *Research in Astronomy and Astrophysics* 20.11, 188, p. 188. DOI: 10.1088/1674-4527/20/11/188. arXiv: 2007.01015 [astro-ph.HE].

Xu, Kun and Xiang-Dong Li (June 2019). "On the Fallback Disk around the Slowest Isolated Pulsar, 1E 161348-5055". In: *ApJ* 877.2, 138, p. 138. DOI: 10.3847/1538-4357/ab1902. arXiv: 1906.07857 [astro-ph.HE].

Xu, Kun et al. (Aug. 2021). "Do the Periodic Activities of Repeating Fast Radio Bursts Represent the Spins of Neutron Stars?" In: *ApJ* 917.1, 2, p. 2. DOI: 10.3847/1538-4357/ac05ba. arXiv: 2105.13122 [astro-ph.HE].

Xu, Kun et al. (Apr. 2023). "Back to the Starting Point: On the Simulation of the Initial Magnetic Fields and Spin Periods of Nonaccreting Pulsars". In: *ApJ* 947.2, 76, p. 76. DOI: 10.3847/1538-4357/acc8ce. arXiv: 2304.03530 [astro-ph.HE].

Yadan, Omry (2019). *Hydra - A framework for elegantly configuring complex applications*. Github. URL: https://github.com/facebookresearch/hydra.

Yan, T., R. Perna, and R. Soria (July 2012). "The influence of fallback discs on the spectral and timing properties of neutron stars". In: *MNRAS* 423.3, pp. 2451–2463. DOI: 10.1111/j.1365-2966.2012.21051.x. arXiv: 1205.3259 [astro-ph.HE].

Yao, J. M., R. N. Manchester, and N. Wang (Jan. 2017). "A New Electron-density Model for Estimation of Pulsar and FRB Distances". In: *ApJ* 835.1, 29, p. 29. DOI: 10.3847/1538-4357/835/1/29. arXiv: 1610.09448 [astro-ph.GA].

Yusifov, I. and I. Küçük (2004). "Revisiting the radial distribution of pulsars in the Galaxy". In: *A&A* 422, pp. 545–553. DOI: 10.1051/0004-6361:20040152. arXiv: astro-ph/0405559 [astro-ph].

Zanni, C. and J. Ferreira (Feb. 2013). "MHD simulations of accretion onto a dipolar magnetosphere. II. Magnetospheric ejections and stellar spin-down". In: *A&A* 550, A99, A99. DOI: 10.1051/0004-6361/201220168. arXiv: 1211.4844 [astro-ph.SR].

Zhang, Bing (July 2023). "The physics of fast radio bursts". In: *Reviews of Modern Physics* 95.3, 035005, p. 035005. DOI: 10.1103/RevModPhys.95.035005. arXiv: 2212.03972 [astro-ph.HE].

Zhang, Bing and Janusz Gil (Oct. 2005). "GCRT J1745-3009 as a Transient White Dwarf Pulsar". In: *ApJ* 631.2, pp. L143–L146. DOI: 10.1086/497428. arXiv: astro-ph/0508213 [astro-ph].

Zhang, Bing, Alice K. Harding, and Alexander G. Muslimov (Mar. 2000). "Radio Pulsar Death Line Revisited: Is PSR J2144-3933 Anomalous?" In: *ApJ* 531.2, pp. L135–L138. DOI: 10.1086/312542. arXiv: astro-ph/0001341 [astro-ph].

Zhong, Yici et al. (Aug. 2021). "A Necessary Condition for Supernova Fallback Invading Newborn Neutron-star Magnetosphere". In: *ApJ* 917.2, 71, p. 71. DOI: 10.3847/1538-4357/ac0a74. arXiv: 2103.09461 [astro-ph.HE].