

Analyzing the Galactic pulsar distribution with machine learning

M. RONCHI ^{1,2} V. GRABER ^{1,2} A. GARCIA-GARCIA ^{1,2,3} N. REA ^{1,2} AND J. A. PONS ³

¹*Institute of Space Sciences (ICE, CSIC), Campus UAB, Carrer de Can Magrans s/n, 08193, Barcelona, Spain*

²*Institut d'Estudis Espacials de Catalunya (IEEC), Carrer Gran Capità 2-4, 08034 Barcelona, Spain*

³*Departament de Física Aplicada, Universitat d'Alacant, 03690 Alicante, Spain*

(Accepted ApJ)

ABSTRACT

We explore the possibility of inferring the properties of the Galactic neutron star population through machine learning. In particular, in this paper we focus on their dynamical characteristics and show that an artificial neural network is able to estimate with high accuracy the parameters which control the current positions of a mock population of pulsars. For this purpose, we implement a simplified population-synthesis framework (where selection biases are neglected at this stage) and concentrate on the natal kick-velocity distribution and the distribution of birth distances from the Galactic plane. By varying these and evolving the pulsar trajectories in time, we generate a series of simulations that are used to train and validate a suitably structured convolutional neural network. We demonstrate that our network is able to recover the parameters governing the kick-velocity and Galactic height distribution with a mean relative error of about 10^{-2} . We discuss the limitations of our idealized approach and study a toy problem to introduce selection effects in a phenomenological way by incorporating the observed proper motions of 216 isolated pulsars. Our analysis highlights that increasing the sample of pulsars with accurate proper motion measurements by a factor of ~ 10 , one of the future breakthroughs of the Square Kilometer Array, we might succeed in constraining the birth spatial and kick-velocity distribution of the neutron stars in the Milky Way with high precision through machine learning.

Keywords: stars: neutron – pulsars: general – proper motions – neural networks: convolutional

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 & 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 & Kramer 2004; 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 & McCammon 1992; Taylor & Cordes 1993; Cordes & Lazio 2002; Yao et al. 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 et al. 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 & Phinney 1997; Lai et al. 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 & Cordes 1987; Mandel & Müller 2020). Additionally, the anisotropic emission of neutrinos has been suggested to impart kicks on compact remnants (Bisnovatyi-Kogan 1993; Fryer & Kusenko 2006; Tamborra et al. 2014; Nagakura et al. 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). 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 et al. 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 et al. 2002; Brisken et al. 2003; Hobbs et al. 2005; Faucher-Giguère & Kaspi 2006; Gullón et al. 2014; Verbunt et al. 2017; Cieřlar et al. 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.

Constraining the birth properties of isolated neutron stars from observational data is the main motivator for our work. Instead of following earlier approaches that have employed standard statistical techniques, we focus on characterizing initial pulsar properties using machine learning techniques which have seen increasing interest in the astronomy and astrophysics communities in recent years. This paper, which will be the first in a series, is dedicated to the technical aspects of these efforts and aims to 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, i.e., with the Square Kilometer Array (SKA).

Our paper is structured as follows: In §2 we describe the methods used to simulate and evolve a mock neutron star population in time. §3 contains a description of the machine-learning framework, including the generation of our data-sets (§3.1), the employed network architectures (§3.2) as well as details of the training process (§3.3). In §4 we present our experiments, which are discussed in detail and connected with observational data in §5. Finally, we provide a summary and outlook in §6.

2. POPULATION SYNTHESIS

A widely used approach to investigate the properties of the observed neutron star population is through population synthesis (see e.g. Narayan & Ostriker 1990; Faucher-Giguère & Kaspi 2006; Gonthier et al. 2007; Kiel et al. 2008; Kiel & Hurley 2009; Osłowski et al. 2011; Levin et al. 2013; Gullón et al. 2014; Bates et al. 2014; Cieřlar et al. 2020). 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 & Kaspi (2006), Gullón et al. (2014) and Cieřlar et al. (2020). The necessary ingredients are briefly summarized in the following.

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 et al. 2021). This yields a total of 10^5 simulated neutron stars for each synthetic population, whose evolution we can compute within reasonable timescales.

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 center of the Galaxy. Here r represents the distance in kiloparsec from the Galactic center, 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:

$$\begin{cases} x = r \cos \phi, \\ y = r \sin \phi, \\ z = z. \end{cases} \quad (1)$$

We assume that the Sun is located at the coordinates $x = 0$ kpc, $y = R_{\odot}$, $z = z_{\odot}$, where $R_{\odot} = 8.3$ kpc and $z_{\odot} = 0.02$ 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 center 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 Eq. (15) in Yusifov & Küçük (2004) to determine the probability density function for the radial distance:

$$P(r) = 2\pi r \rho(r), \quad (2)$$

with

$$\rho(r) = A \left(\frac{r + R_1}{R'_{\odot} + R_1} \right)^a \exp \left[-b \left(\frac{r - R'_{\odot}}{R'_{\odot} + R_1} \right) \right], \quad (3)$$

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 center. Although different from the R_{\odot} value assumed above, we keep $R'_{\odot} = 8.5$ kpc in this parameterization in order to be consistent with the results of Yusifov & Küçük (2004). We note that this is the distribution for evolved pulsars rather than that of their progenitors and Yusifov & Küçük (2004) find small discrepancies between this distribution and that of OB stars. However, Faucher-Giguère & Kaspi (2006) show that the evolved pulsar population is well described by birth positions drawn from Eq. (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 center:

$$\phi(r) = k \ln \left(\frac{r}{r_0} \right) + \phi_0. \quad (4)$$

Our values of the model parameters, i.e., the winding constant k , the inner radius r_0 and the inner

Table 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 et al. (2017); see §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

angle ϕ_0 are reported in Table 1 and evaluated from Table 1 in Yao et al. (2017) in order to match the same functional form as defined in Eq. (4). For our analysis, we follow Faucher-Giguère & Kaspi (2006) and do not include the Local arm, whose density is much smaller than that of the four major arms (Cordes & Lazio 2002; Yao et al. 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 Eq. (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 birth position, we add noise to both coordinates r and ϕ to smear out the distribution and avoid artificial features near the Galactic center. For this purpose we add a correction $\phi_{\text{corr}} = \phi_{\text{rand}} \exp(-0.35r/\text{kpc})$ 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 centered at 0 with standard deviation $\sigma = 0.07r$. Although this prescription was introduced by Faucher-Giguère & 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 $(r + r_{\text{corr}}, \phi(r) + \Omega t_{\text{age}} + \phi_{\text{corr}})$ 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 :

$$P(z) = \frac{1}{h_c} \exp\left(-\frac{|z|}{h_c}\right). \quad (5)$$

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 $h_c = 0.18$ 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 $h_c = 0.18$ 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.

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. We consider a progenitor circular orbital speed given by the following relation:

$$v_{\text{orb}} = \sqrt{r \frac{\partial \Phi_{\text{MW}}(r, z)}{\partial r}}, \quad (6)$$

where Φ_{MW} is the Milky Way gravitational potential discussed below. We assume that each neutron star has an initial kick velocity \mathbf{v}_k , whose 3D magnitude v_k is

Table 2. Parameters of the Milky Way gravitational potential taken from Table 1 in [Marchetti et al. \(2019\)](#), see also [Bovy \(2015\)](#).

Component	Parameters
nucleus (n)	$M_n = 1.71 \times 10^9 M_\odot$ $r_n = 0.07$ kpc
bulge (b)	$M_b = 5.0 \times 10^9 M_\odot$ $r_b = 1.0$ kpc
disk (d)	$M_d = 6.8 \times 10^{10} M_\odot$ $a_d = 3.00$ kpc $b_d = 0.28$ kpc
halo (h)	$M_h = 5.4 \times 10^{11} M_\odot$ $r_h = 15.62$ kpc

randomly drawn from a Maxwell distribution, shaped by the dispersion parameter σ_k :

$$P(v_k) = \sqrt{\frac{2}{\pi}} \frac{v_k^2}{\sigma_k^3} \exp\left(-\frac{v_k^2}{\sigma_k^2}\right). \quad (7)$$

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 & Becker 2006](#); [Pavan et al. 2014](#)). Based on pulsar timing measurements, [Hobbs et al. \(2005\)](#) have suggested that $\sigma_k = 265 \text{ km s}^{-1}$ 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 ($v_{k,r}$, $v_{k,\phi}$, $v_{k,z}$) 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 ($v_{k,r}$, $v_{\text{orb}} + v_{k,\phi}$, $v_{k,z}$).

2.4. Galactic Potential

As typical for spiral galaxies like the Milky Way, we assume an axisymmetric Galactic potential Φ_{MW} ([Carlborg & 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](#)

et al. (2019):

$$\Phi_{\text{MW}} = \Phi_n + \Phi_b + \Phi_d + \Phi_h. \quad (8)$$

The nucleus and the bulge are described by a Hernquist potential ([Hernquist 1990](#)):

$$\Phi_n = -\frac{GM_n}{r_n + r}, \quad (9)$$

$$\Phi_b = -\frac{GM_b}{r_b + r}. \quad (10)$$

The disk has a Miyamoto-Nagai disk potential ([Miyamoto & Nagai 1975](#)):

$$\Phi_d = -\frac{GM_d}{\sqrt{K^2 + r^2}}, \quad (11)$$

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 et al. 1996](#)):

$$\Phi_h = -\frac{GM_h}{r} \ln\left(1 + \frac{r}{r_h}\right). \quad (12)$$

The parameters of this model are reported in Table 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.

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: $\ddot{\mathbf{r}} = -\nabla\Phi_{\text{MW}}$. In cylindrical galactocentric coordinates the three components of this vector equation take the form:

$$\begin{cases} \ddot{r} - r\dot{\phi}^2 = -\frac{\partial\Phi_{\text{MW}}}{\partial r}, \\ 2\dot{r}\dot{\phi} + r\ddot{\phi} = -\frac{1}{r}\frac{\partial\Phi_{\text{MW}}}{\partial\phi}, \\ \ddot{z} = -\frac{\partial\Phi_{\text{MW}}}{\partial z}. \end{cases} \quad (13)$$

For each neutron star we numerically integrate the above equations in time from $t = 0 \text{ yr}$ to $t = t_{\text{age}}$ 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 et al. 2015)¹. To assess 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, 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.

3. MACHINE-LEARNING SET-UP

In the past decade, the accumulation of extensive and heavy data-sets 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 modeling 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 & Brunner 2010; Allen et al. 2019; Baron 2019; Fluke & Jacobs 2020). For compact-object related science, ML algorithms have for example been developed to classify new pulsars candidates (Bethapudi & Desai 2018; Balakrishnan et al. 2020; Lin et al. 2020) as well as transient radio events such as FRBs (Agarwal et al. 2020). Other approaches have aimed at forecasting and analyzing gravitational-wave signals in real time (Cabero et al. 2020; Wei & Huerta 2020; Skliris et al. 2020; Gerosa et al. 2020), interpreting gravitational wave-events in light of population synthesis (Wong & Gerosa 2019), or reconstructing the neutron star equation of state from observed quantities (Morawski & Bejger 2020).

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. 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 organized 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 data-set 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 minimize a specific loss function which measures the distance between the network output prediction and the target ground truth.

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 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, then (ii) focus on the specific network architecture employed, and finally (iii) describe the details of our training process.

3.1. Data-set 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 §2, we perform the following simulation runs:

S1 We generate 10 data-sets 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

¹ `Numba` translates Python functions into optimized machine code at run-time, which allows us to achieve a speed-up by about a factor of 6.




Figure 1. Examples of 128×128 resolution maps in the galactocentric xy -plane extending from -20 to 20 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 $h_c = 0.18$ kpc and $\sigma_k = 265$ km s $^{-1}$. For visualization purposes, we represent the data using a colormap to highlight the resulting structures; red regions are characterized 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 recognizable 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 behavior because the spiral-arm structure is smeared out.




Figure 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 visualization purposes, we represent the data using a colormap to highlight the resulting structures; red regions are characterized 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 Center. 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.

velocity distribution in the range $[1, 700]$ km s $^{-1}$.² We also generate a test data-set 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 data-set 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 data-set 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 data-sets, where we uniformly vary the kick-velocity parameter σ_k parameter 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 data-set 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 data-set 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.

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.

As addressed in detail in §4, the smaller simulation data-sets will be used to explore the network behavior. The largest data-sets containing 20000 and 16384 samples, respectively, and the test data-sets with 1000 samples will be used to perform the final training experiments and assess the actual network accuracy in generalization 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 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 $4.0 \cdot \sigma + 0.5$ 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 generalization 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 $[\text{km s}^{-1}]$.
- 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 $[\text{mas yr}^{-1}]$.

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 data-sets 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 $h_c = 0.18$ kpc and $\sigma_k = 265 \text{ km s}^{-1}$ is shown in Figs. 1 and 2.

Before loading the maps into our ML pipeline, they are normalized 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 normalization is to speed up the training process and make convergence easier since all inputs will provide signals of similar magnitude to the loss-function minimization. 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 normalization might lead to slower, worse or even no convergence at all. Apart from the blurring and normalization described above, we do not apply any additional pre-processing steps to our input data.

3.2. Network Architecture

For the implementation of our ML pipeline we use PyTorch (Paszke et al. 2019), an optimized 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 connected neural network also referred to as a multi-layer perceptron (MLP) 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 $\text{ReLU}(x) = \max(0, x)$. 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 classifi-




Figure 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.

cation tasks when applied to structured and matrix-like 2D inputs (see Rawat & 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 3. A schematic representation of its structure for a 4-channel input with galactocentric maps is also shown in Fig. 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.
- 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).

Table 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)

For the convolutional and pooling layers the stride parameter regulates the amount of displacement in bins that 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.

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 data-set different from the training set. This is needed to assess the ability of the network to generalize what it is learning to an unknown data-set. The minimization 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 initialize the weights and biases. For this purpose we use the Kaiming initialization 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 optimization step is performed. In general, for bigger training data-sets, 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. For the gradient-descent optimizer we use the Adaptive Moment Estimation (Adam) (Kingma & 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. When using the Adam optimizer 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 optimization steps by going through the whole training data-set 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 optimized 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 generalize over unknown data.

4. EXPERIMENTS

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 neu-

ral 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 behavior 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 generalization power. A similar strategy is then followed for the two-parameter prediction.

4.1. Single-parameter Predictions

4.1.1. 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 data-set with 20000 samples as it will be used to assess the generalization power of the best performing network (see §4.1.2). Thus, we have training sets 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 §3.2 and §3.3. For these initial experiments we do not incorporate validation but only focus on the training results for different types of data-sets, as we are not yet interested in evaluating the generalization 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 km s^{-1} . If the final RMSE value is higher than this threshold the training is repeated up

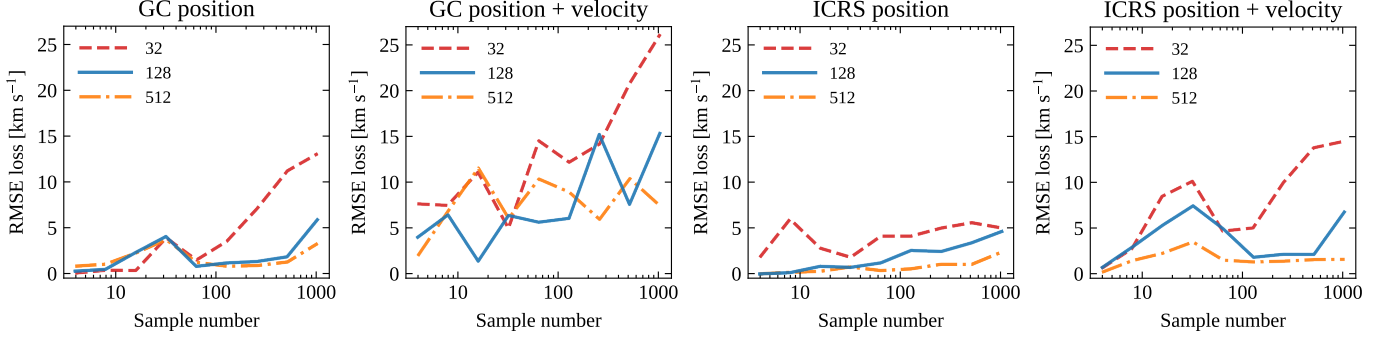


Figure 4. MLP best training RMSE values for the training process on the single parameter σ_k of the Maxwell kick-velocity distribution, as a function of the training data-set size and the resolution (*red*, *blue*, and *orange* curves for 32, 128 and 512 respectively) using the four different input configurations T1 (GC position), T2 (GC position + velocity), T3 (ICRS position) and T4 (ICRS position + velocity).

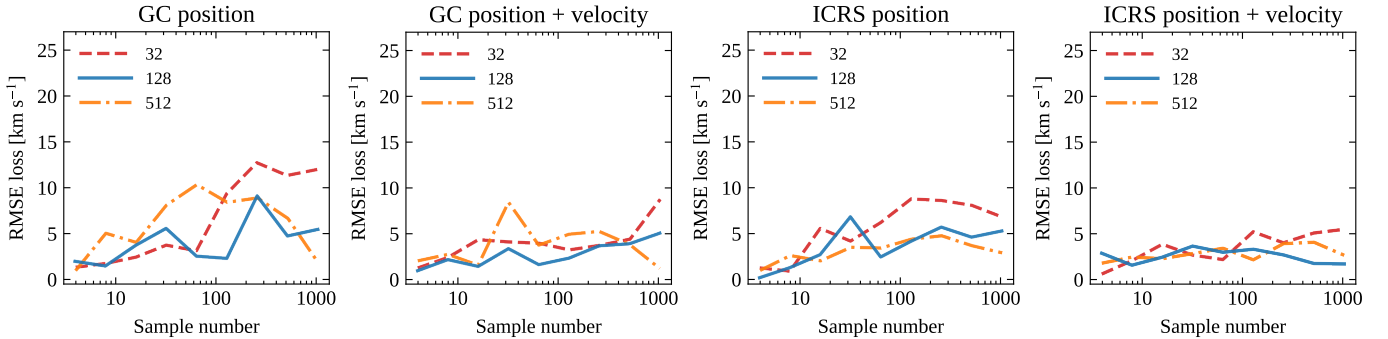


Figure 5. CNN best training RMSE values for the training process on the single parameter σ_k of the Maxwell kick-velocity distribution, as a function of the training data-set size and the resolution (*red*, *blue*, and *orange* curves for 32, 128 and 512 respectively) using the four different input configurations T1 (GC position), T2 (GC position + velocity), T3 (ICRS position) and T4 (ICRS position + velocity).

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 optimization step on a single data batch (see Appendix B 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 data-set sizes 4, 8, 16, 32 and progressively

increase it to 4, 8, 16, 32, 64 as the data-set 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 data-sets from S1 with configuration T1, T2, T3 and T4 are shown in Figs. 4 and 5, respectively (see Appendix B for the timing results), highlighting the best converged RMSE values as a function of the training data-set 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

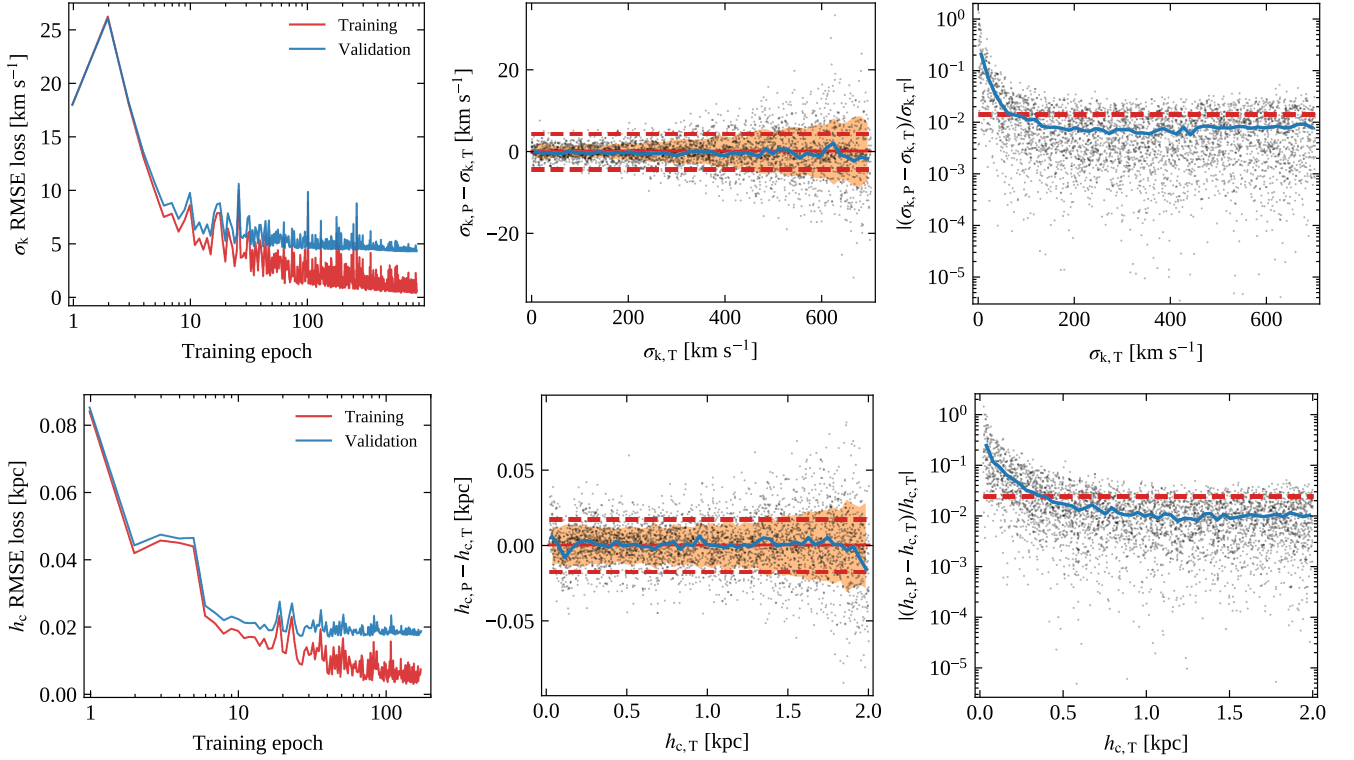


Figure 6. Results of the CNN single-parameter prediction for the kick-velocity parameter σ_k and scale height h_c for the corresponding validation sets. *Top (bottom) left panel:* evolution of the RMSE training (red) and validation (blue) losses as a function of the training/validation epoch for the 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 km s^{-1} (0.017 kpc) computed over the whole range $[1, 700] \text{ km s}^{-1}$ ($[0.02, 2] \text{ kpc}$). 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 centered 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 $[1, 700] \text{ km s}^{-1}$ ($[0.02, 2] \text{ kpc}$). 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.

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 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 Figs. 13 and 14 in Appendix B). 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 em-

ploy our CNN architecture in the following experiments, which should also be less prone to overfitting.

4.1.2. Generalization Results

As the next step we separately train the CNN to predict the σ_k and h_c parameters by using the two big data-sets 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 data-sets 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 data-sets, 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. The evolution of the training and validation losses is shown in the left panels of Fig. 6.

The predictions of the trained network on the validation set for σ_k and h_c are summarized in Fig. 6 and Table 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 data-set with a RMSE uncertainty of 4.4 km s^{-1} , computed over the whole range $[1, 700] \text{ km s}^{-1}$. This is indicated by the red dashed lines in the residuals plot (see top central panel of Fig. 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 data-set with a RMSE uncertainty of 0.017 kpc, computed over the whole range $[0.02, 2] \text{ kpc}$ (see bottom central panel of Fig. 6). Note that in both experiments, the residuals spread out as the target values increase. We visualize 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 Fig. 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 stabilizes 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 generalization 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 generalize well over unseen data-sets.

4.2. Two-parameter Predictions

4.2.1. 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 data-set of maps generated from simulation run S3 (see §3.1). In this case, we have training sets with increasing number of samples, $16 = 4 \times 4$, $64 = 8 \times 8$, $256 = 16 \times 16$, $1024 = 32 \times 32$, $4096 = 64 \times 64$. We leave aside the largest data-set with 16384 = 128x128 simulations for our final generalization 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 §4.1.1). We fix the CNN architecture and the training set-up as described in §3.2 and §3.3, respectively. As for the single-parameter comparison, we only focus on the training behavior for this initial comparison. We keep track of the RMSE training losses for both parameters separately, but training proceeds by minimizing 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 optimization step (see Appendix B). The initial learning rate is set to 10^{-4} , while the batch size is changed according to the data-set size. In particular, we use a batch size of 1, 4, 16, 64,

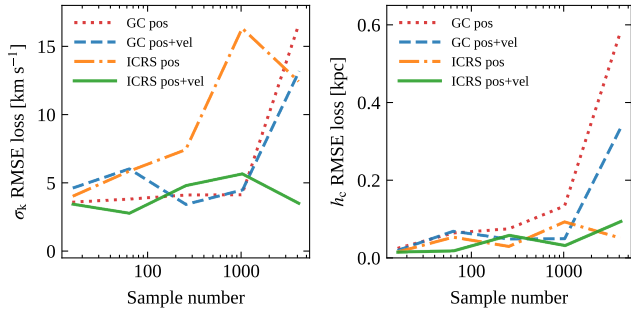


Figure 7. CNN best training RMSE values for the training process on the two parameters σ_k of the Maxwell kick-velocity distribution (*left panel*) and characteristic scale height h_c (*right panel*) as a function of the training data-set size for the four different input configurations T1 (GC position), T2 (GC position + velocity), T3 (ICRS position) and T4 (ICRS position+ velocity).

128 for the data-set sizes 16, 64, 256, 1024 and 4096, respectively.

The results of our experiments using the CNN for the two-parameter prediction on the training data-sets from S3, with configuration T1, T2, T3 and T4 are summarized in Fig. 7 (see Appendix B 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 $\lesssim 5 \text{ km s}^{-1}$ for the σ_k parameter, comparable with the single-parameter case. For h_c , the CNN reaches a training RMSE $\lesssim 0.1 \text{ kpc}$. However, we note a drop in accuracy for the σ_k parameter when only the ICRS density maps are used. As already mentioned in §4.1.1, 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. 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.

4.2.2. Generalization 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 data-set 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 data-set 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 data-set and use the same configuration as in §4.1.2. The evolution of the individual training and validation losses is shown in the left panels of Fig. 8.

The results for the trained network’s prediction on the validation set for both parameters are shown in Fig. 8 and summarized in Table 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 Fig. 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 Fig. 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 Fig. 8.

We then evaluate the generalization 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 generalization 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

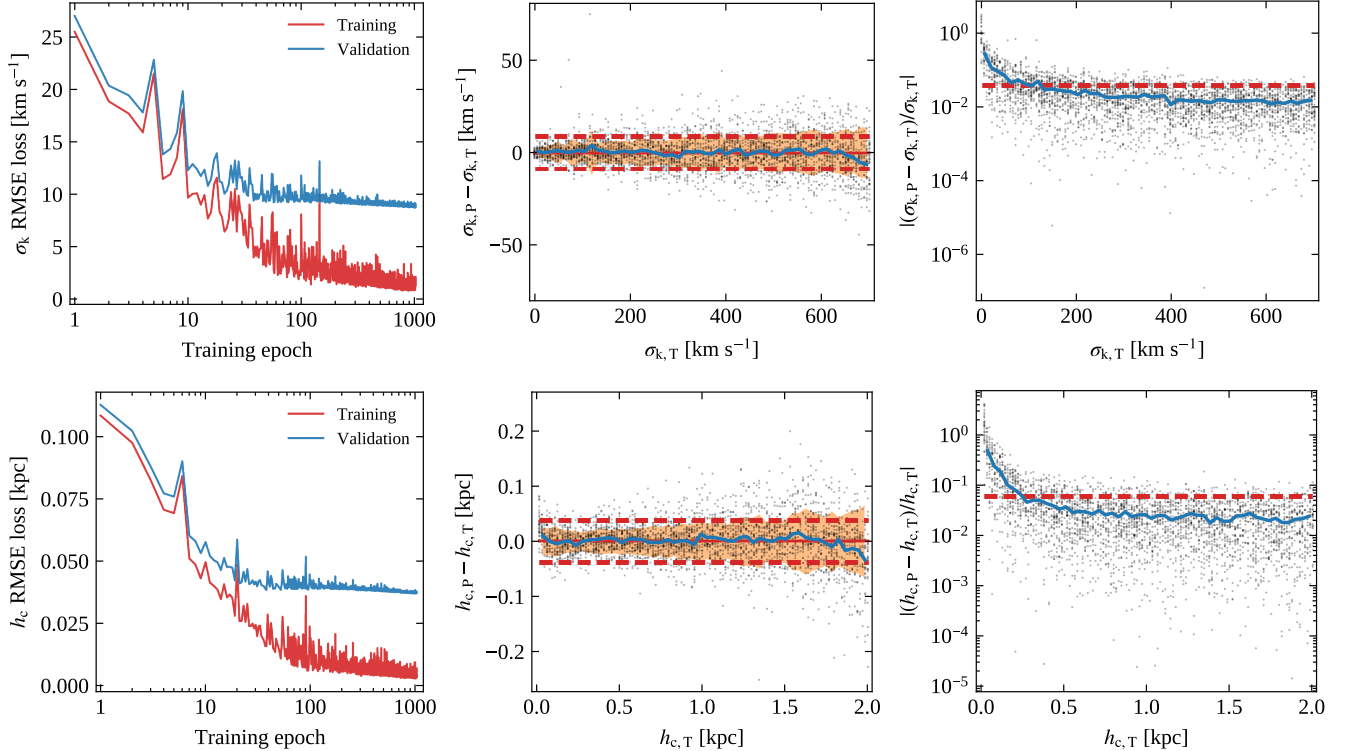


Figure 8. Results of the CNN two-parameter prediction for the kick-velocity parameter σ_k and scale height h_c for the corresponding validation sets. *Top (bottom) left panel:* evolution of the RMSE training (red) and validation (blue) losses as a function of the training/validation epoch for the σ_k (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 $[1, 700] \text{ km s}^{-1}$ ($[0.02, 2] \text{ kpc}$). 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 centered 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 $[1, 700] \text{ km s}^{-1}$ ($[0.02, 2] \text{ kpc}$). 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.

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 data-set used above where both σ_k and h_c are varied. After predicting on the validation data-set 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: Fig. 9 shows the residuals of the scale-height parameter h_c versus the residuals of the kick-velocity parameter σ_k for the predictions 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. DISCUSSION

In this paper 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-

Table 4. Summary of the CNN generalization results on the validation and test (in parenthesis) data-sets for the single-parameter and two-parameter cases.

Parameter	1-par. generalization		2-par. generalization	
	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)

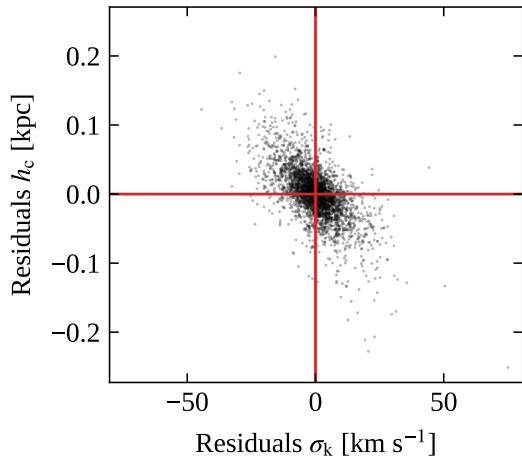


Figure 9. Scatter plot of residuals of the predicted scale-height parameter h_c versus residuals of the predicted kick-velocity parameter σ_k for the validation data-set of our two-parameter generalization experiment. An anticorrelation can be observed.

velocity distribution and the distribution of birth distances from the Galactic plane. Taking into account the Galaxy 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 generalized results presented in the previous sections are obtained in a very idealized 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 et al. 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 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; Cieřlar et al. 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 et al. (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 & Kaspi (2006) studied the velocity component along the Galactic longitude for a sample of 34 pulsars observed with interferometric techniques (Brisken et al. 2002, 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 et al. (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 & Müller 2020). While also possible scenarios for isolated systems (Janka 2017), such conditions might generally affect those neutron stars born in binaries (Giacobbo & Mapelli 2020), where mass-loss episodes could strip their progenitors off their hydrogen




Figure 10. Proper motions of the selected 216 neutron stars in the sky represented in the ICRS reference frame. The current locations of the neutron stars are indicated by the colored circles, whereas the tracks indicate their motion for the past 0.5 Myr, assuming no radial velocity and neglecting the effects of the Galactic potential. The color encodes the heliocentric distance d_{\odot} of the neutron stars. The corresponding data is provided in Table 5. In the background, we show in gray 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 Center.

envelopes. This might favor the formation of small iron cores or accretion-induced electron-capture supernovae, resulting in weaker natal kicks (Schwab et al. 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 model but could in principle generate an imprint on the observed neutron star population. The ability of ML algorithms to recognize features and patterns in the processed data could help to better constrain the bi-modality of the kick-velocity distribution. In future work, we will investigate these aspects in more detail and explore ways to implement a ML framework to reconstruct the shape of the kick-velocity distribution underlying an evolved population of pulsars.

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 & Manchester 1998; Melrose 2017), 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). 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 could potentially influence a future comparison with observational 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 catalogue³ (Manchester et al. 2005), but in some cases we provide proper motion results from more recent analyses (see Appendix C 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 5. 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 DM s using the free-electron density model of Yao et al. (2017)

³ <https://www.atnf.csiro.au/research/pulsar/psrcat/>

(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 Fig. 10.

In the left panel of Fig. 11, the gray 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 Center (see the red histogram in the left panel of Fig. 11). Thus, the shape of the gray distribution in Fig. 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 Fig. 11, we also note that a selection bias on distance is also reflected in the distribution of the total proper motion magnitudes (gray histogram), computed as $\mu_{\text{tot}} \equiv \sqrt{\mu_{\text{RA}}^2 + \mu_{\text{DEC}}^2}$. Indeed, since the nearest stars are also characterized by larger angular proper motions, there is a tendency to detect high proper-motion stars with higher probability.

In this first empirical approach, instead of attempting to identify these underlying biases (which will be the subject of a subsequent paper in preparation), 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 $\sigma_k = 265 \text{ km s}^{-1}$ and $h_c = 0.18 \text{ kpc}$, respectively. 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 Fig. 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 underlying 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 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

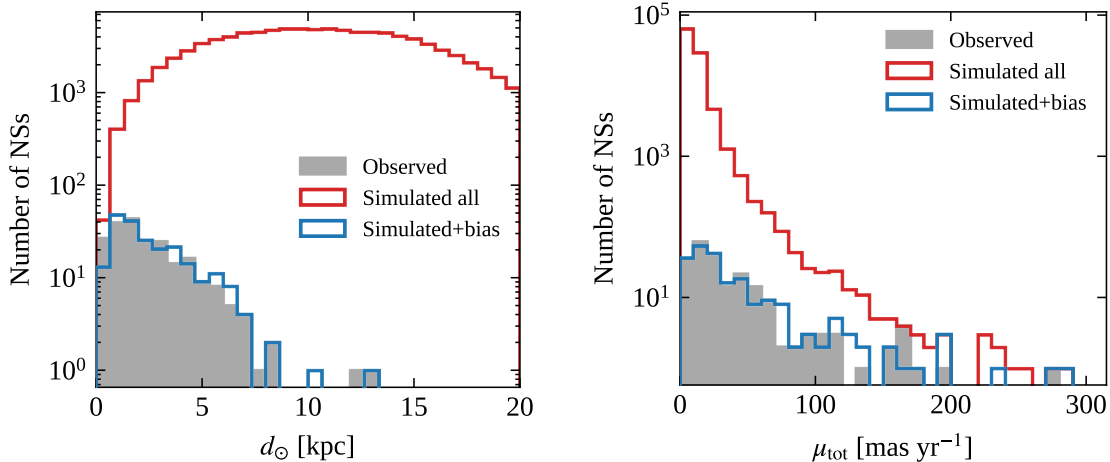


Figure 11. Distribution of heliocentric distances d_{\odot} (left panel) and proper-motion magnitudes μ_{tot} (right panel) for the 216 neutron stars with measured proper motion (gray histograms). For comparison, we also show the distances and proper-motion magnitudes for 216 simulated neutron stars (blue histograms) randomly sampled from a mock population generated with the fiducial parameters $\sigma_k = 265 \text{ km s}^{-1}$ and $h_c = 0.18 \text{ kpc}$ (red histograms). For the weighted sampling, we use the weight function $f(d_{\odot}) = d_{\odot}^{-1} \exp(-0.5d_{\odot})$.

same training hyperparameters as in §4.1.2, 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 Fig. 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.

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 characterizing 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, 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 data-set 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).

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 pre-




Figure 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 $f(d_\odot) = d_\odot^{-1} \exp(-0.5d_\odot)$.

dicting σ_k and h_c from ICRS maps, the network is able to reach absolute uncertainties lower than 10 km s^{-1} and 0.05 kpc , respectively, corresponding to a relative error of around 10^{-2} for both parameters. Albeit obtained under simplified assumptions, our feasibility study, the main focus of this paper, 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 modeling 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. The potential to learn complex patterns from multi-dimensional data and the power of generalizing to unseen data-sets, makes ML algorithms a powerful tool to tackle such multi-parameter optimization studies.

ACKNOWLEDGEMENTS

We would like to thank Alice Borghese and Francesco Coti Zelati for their help with the collection of the observed pulsar proper motions, Daniele Viganò for providing comments on the manuscript, and Alessandro Patruno for helpful conversations related to this paper. MR, VG, AG and NR acknowledge support from the H2020 ERC Consolidator Grant “MAGNESIA” under grant agreement Nr. 817661 (PI: Rea), and grants SGR2017-1383 and PGC2018-095512-B-I00. JAP acknowledges support from the Generalitat Valenciana grant PROMETEO/2019/071, AEI grant PGC2018-095984-B-I00 and the Alexander von Humboldt Foundation via a Humboldt Research Award. This work has also been partially supported by the PHAROS COST Action (CA16214). The data production, processing and analysis tools for this paper have been developed, implemented and operated in collaboration with the Port d’Informació Científica (PIC) data center. PIC is maintained through a consortium of the Institut de Física d’Altes Energies (IFAE) and the Cen-

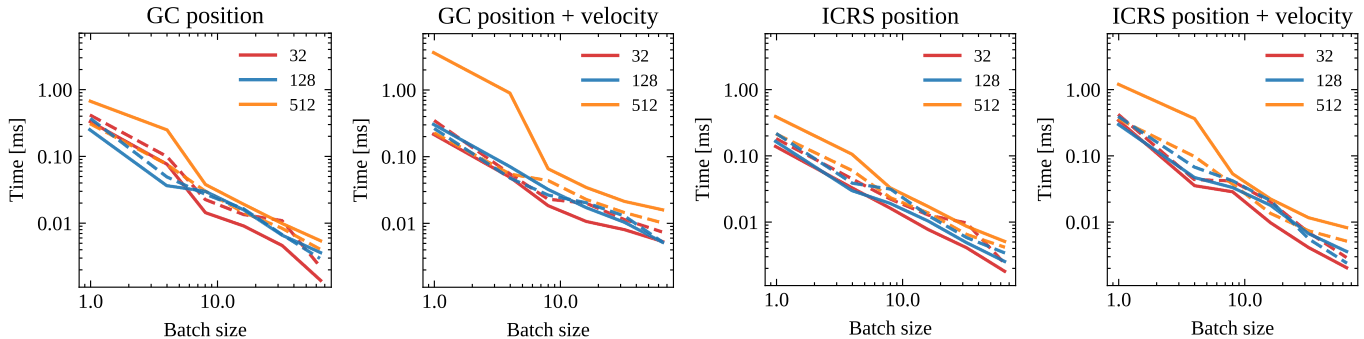


Figure 13. MLP forward (*solid lines*) and backward (*dashed lines*) pass times per sample in ms for the training process on the single parameter σ_k of the Maxwell kick-velocity distribution, as a function of the batch size and the resolution (*red, blue, and orange* curves for 32, 128 and 512 respectively) using the four different input configurations T1 (GC position), T2 (GC position + velocity), T3 (ICRS position) and T4 (ICRS position + velocity).

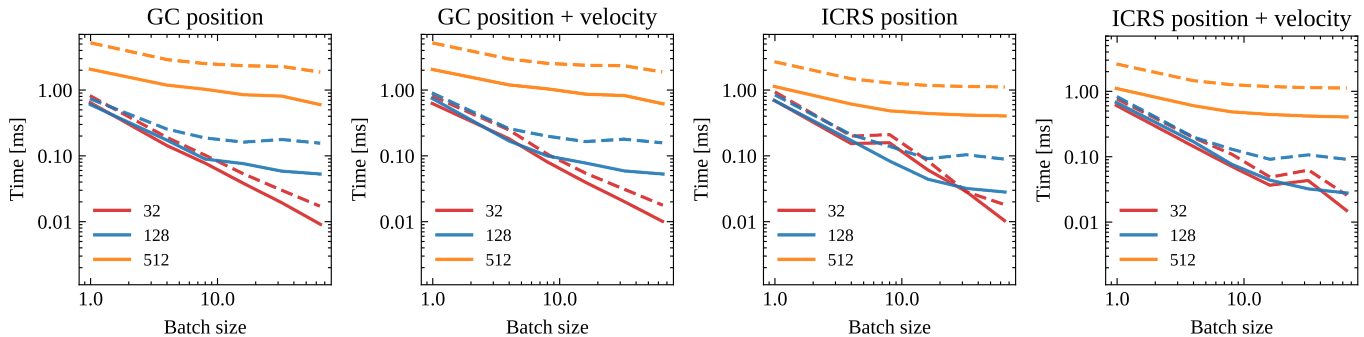


Figure 14. CNN forward (*solid lines*) and backward (*dashed lines*) pass times per sample in ms for the training process on the single parameter σ_k of the Maxwell kick-velocity distribution, as a function of the batch size and the resolution (*red, blue, and orange* curves for 32, 128 and 512 respectively) using the four different input configurations T1 (GC position), T2 (GC position + velocity), T3 (ICRS position) and T4 (ICRS position + velocity).

tro de Investigaciones Energéticas, Medioambientales y Tecnológicas (Ciemat). We particularly thank Christian Neissner, Ricard Cruz, Carles Acosta, Gonzalo Merino, and Pau Tallada for their support at PIC. Finally, we acknowledge the use of the following software: *Astropy* (Astropy Collaboration et al. 2013, 2018),

Hydra (Yadan 2019), *IPython* (Perez & Granger 2007), *Matplotlib* (Hunter 2007), *Numba* (Lam et al. 2015), *NumPy* (Oliphant 2006; van der Walt et al. 2011; Harris et al. 2020), *Pandas* (McKinney 2010), *PyTorch* (Paszke et al. 2019), and *SciPy* (Jones et al. 2001; Virtanen et al. 2020).

APPENDIX

A. 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.

B. TIMING TESTS

B.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 §4.1.1. 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 optimization step) as a function of the batch size and resolution. Our

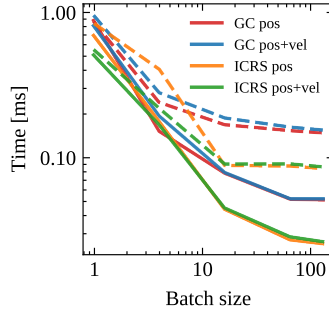


Figure 15. CNN forward pass (*solid line*) and backward pass time (*dashed line*) per sample for the two-parameter experiments as a function of the batch size for the four different input configurations T1 (GC position), T2 (GC position + velocity), T3 (ICRS position) and T4 (ICRS position + velocity).

benchmarks for the training data-sets from simulation run S1 using the four training configurations T1, T2, T3 and T4 are shown in Figs. 13 and 14, respectively.

The timing benchmark shows that the MLP is slightly faster in performing an optimization 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 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).

B.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 Fig. 15. 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.

C. NEUTRON STARS WITH MEASURED PROPER MOTION

In Table 5, 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 catalogue⁴ (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 [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

⁴ <https://www.atnf.csiro.au/research/pulsar/psrcat/>

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).

REFERENCES

- Agarwal, D., Aggarwal, K., Burke-Spolaor, S., Lorimer, D. R., & Garver-Daniels, N. 2020, MNRAS, 497, 1661, doi: [10.1093/mnras/staa1856](https://doi.org/10.1093/mnras/staa1856)
- Allen, G., Andreoni, I., Bachelet, E., et al. 2019, preprint, arXiv:1902.00522. <https://arxiv.org/abs/1902.00522>
- Arzoumanian, Z., Chernoff, D. F., & Cordes, J. M. 2002, ApJ, 568, 289, doi: [10.1086/338805](https://doi.org/10.1086/338805)
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, A&A, 558, A33, doi: [10.1051/0004-6361/201322068](https://doi.org/10.1051/0004-6361/201322068)
- Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, AJ, 156, 123, doi: [10.3847/1538-3881/aabc4f](https://doi.org/10.3847/1538-3881/aabc4f)
- Balakrishnan, V., Champion, D., & Barr, E. 2020, preprint, arXiv:2010.07457. <https://arxiv.org/abs/2010.07457>
- Ball, N. M., & Brunner, R. J. 2010, Int. J. Mod. Phys. D, 19, 1049, doi: [10.1142/S0218271810017160](https://doi.org/10.1142/S0218271810017160)
- Balucinska-Church, M., & McCammon, D. 1992, ApJ, 400, 699, doi: [10.1086/172032](https://doi.org/10.1086/172032)
- Baron, D. 2019, preprint, arXiv:1904.07248. <https://arxiv.org/abs/1904.07248>
- Bates, S. D., Lorimer, D. R., Rane, A., & Swiggum, J. 2014, MNRAS, 439, 2893, doi: [10.1093/mnras/stu157](https://doi.org/10.1093/mnras/stu157)
- Bethapudi, S., & Desai, S. 2018, Astronomy and Computing, 23, 15, doi: [10.1016/j.ascom.2018.02.002](https://doi.org/10.1016/j.ascom.2018.02.002)
- Bisnovatyi-Kogan, G. S. 1993, Astronomical and Astrophysical Transactions, 3, 287, doi: [10.1080/10556799308230566](https://doi.org/10.1080/10556799308230566)
- Bovy, J. 2015, ApJS, 216, 29, doi: [10.1088/0067-0049/216/2/29](https://doi.org/10.1088/0067-0049/216/2/29)
- Brisken, W. F., Benson, J. M., Goss, W. M., & Thorsett, S. E. 2002, ApJ, 571, 906, doi: [10.1086/340098](https://doi.org/10.1086/340098)
- Brisken, W. F., Fruchter, A. S., Goss, W. M., Herrnstein, R. M., & Thorsett, S. E. 2003, AJ, 126, 3090, doi: [10.1086/379559](https://doi.org/10.1086/379559)
- Cabero, M., Mahabal, A., & McIver, J. 2020, ApJL, 904, L9, doi: [10.3847/2041-8213/abc5b5](https://doi.org/10.3847/2041-8213/abc5b5)
- Carlberg, R. G., & Innanen, K. A. 1987, AJ, 94, 666, doi: [10.1086/114503](https://doi.org/10.1086/114503)
- Chatterjee, S., Vlemmings, W. H. T., Brisken, W. F., et al. 2005, ApJL, 630, L61, doi: [10.1086/491701](https://doi.org/10.1086/491701)
- Chen, B. Q., Huang, Y., Hou, L. G., et al. 2019, MNRAS, 487, 1400, doi: [10.1093/mnras/stz1357](https://doi.org/10.1093/mnras/stz1357)
- Cieslar, M., Bulik, T., & Osłowski, S. 2020, MNRAS, 492, 4043, doi: [10.1093/mnras/staa073](https://doi.org/10.1093/mnras/staa073)
- Coenen, T., van Leeuwen, J., Hessels, J. W. T., et al. 2014, A&A, 570, A60, doi: [10.1051/0004-6361/201424495](https://doi.org/10.1051/0004-6361/201424495)
- Cordes, J. M., & Lazio, T. J. W. 2002, preprint, arXiv:astro
- Dang, S. J., Yuan, J. P., Manchester, R. N., et al. 2020, ApJ, 896, 140, doi: [10.3847/1538-4357/ab9082](https://doi.org/10.3847/1538-4357/ab9082)
- Danilenko, A., Karpova, A., Ofengeim, D., Shibanov, Y., & Zyuzin, D. 2020, MNRAS, 493, 1874, doi: [10.1093/mnras/staa287](https://doi.org/10.1093/mnras/staa287)
- Deller, A. T., Tingay, S. J., Bailes, M., & Reynolds, J. E. 2009, ApJ, 701, 1243, doi: [10.1088/0004-637X/701/2/1243](https://doi.org/10.1088/0004-637X/701/2/1243)
- Deller, A. T., Goss, W. M., Brisken, W. F., et al. 2019, ApJ, 875, 100, doi: [10.3847/1538-4357/ab11c7](https://doi.org/10.3847/1538-4357/ab11c7)
- Dewey, R. J., & Cordes, J. M. 1987, ApJ, 321, 780, doi: [10.1086/165671](https://doi.org/10.1086/165671)
- Eisenbeiss, T., Ginski, C., Hohle, M. M., et al. 2010, AN, 331, 243, doi: [10.1002/asna.200911301](https://doi.org/10.1002/asna.200911301)
- Faucher-Giguère, C.-A., & Kaspi, V. M. 2006, ApJ, 643, 332, doi: [10.1086/501516](https://doi.org/10.1086/501516)
- Fluke, C. J., & Jacobs, C. 2020, WIREs Data Mining and Knowledge Discovery, 10, e1349, doi: [10.1002/widm.1349](https://doi.org/10.1002/widm.1349)
- Fryer, C. L., & Kusenko, A. 2006, ApJS, 163, 335, doi: [10.1086/500933](https://doi.org/10.1086/500933)
- Gerosa, D., Pratten, G., & Vecchio, A. 2020, PhRvD, 102, 103020, doi: [10.1103/PhysRevD.102.103020](https://doi.org/10.1103/PhysRevD.102.103020)
- Giacobbo, N., & Mapelli, M. 2020, ApJ, 891, 141, doi: [10.3847/1538-4357/ab7335](https://doi.org/10.3847/1538-4357/ab7335)
- Gonthier, P. L., Story, S. A., Clow, B. D., & Harding, A. K. 2007, Ap&SS, 309, 245, doi: [10.1007/s10509-007-9432-9](https://doi.org/10.1007/s10509-007-9432-9)
- Good, D. C., Andersen, B. C., Chawla, P., et al. 2020, preprint, arXiv:2012.02320. <https://arxiv.org/abs/2012.02320>
- Gullón, M., Miralles, J. A., Viganò, D., & Pons, J. A. 2014, MNRAS, 443, 1891, doi: [10.1093/mnras/stu1253](https://doi.org/10.1093/mnras/stu1253)
- Hansen, B. M. S., & Phinney, E. S. 1997, MNRAS, 291, 569, doi: [10.1093/mnras/291.3.569](https://doi.org/10.1093/mnras/291.3.569)
- Harris, C. R., Jarrod Millman, K., van der Walt, S. J., et al. 2020, Nature, 585, 357, doi: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2)
- He, C., Ng, C. Y., & Kaspi, V. M. 2013, ApJ, 768, 64, doi: [10.1088/0004-637X/768/1/64](https://doi.org/10.1088/0004-637X/768/1/64)

- He, K., Zhang, X., Ren, S., & Sun, J. 2015, preprint, arXiv:1502.01852
- Hernquist, L. 1990, ApJ, 356, 359, doi: [10.1086/168845](https://doi.org/10.1086/168845)
- Hobbs, G., Lorimer, D. R., Lyne, A. G., & Kramer, M. 2005, MNRAS, 360, 974, doi: [10.1111/j.1365-2966.2005.09087.x](https://doi.org/10.1111/j.1365-2966.2005.09087.x)
- Hobbs, G., Lyne, A. G., Kramer, M., Martin, C. E., & Jordan, C. 2004, MNRAS, 353, 1311, doi: [10.1111/j.1365-2966.2004.08157.x](https://doi.org/10.1111/j.1365-2966.2004.08157.x)
- Hui, C. Y., & Becker, W. 2006, A&A, 457, L33, doi: [10.1051/0004-6361:20065941](https://doi.org/10.1051/0004-6361:20065941)
- Hunter, J. D. 2007, Comput. Sci. Eng., 9, 90, doi: [10.1109/MCSE.2007.55](https://doi.org/10.1109/MCSE.2007.55)
- Igoshev, A. P. 2020, MNRAS, 494, 3663, doi: [10.1093/mnras/staa958](https://doi.org/10.1093/mnras/staa958)
- Janka, H.-T. 2017, ApJ, 837, 84, doi: [10.3847/1538-4357/aa618e](https://doi.org/10.3847/1538-4357/aa618e)
- Jennings, R. J., Kaplan, D. L., Chatterjee, S., Cordes, J. M., & Deller, A. T. 2018, ApJ, 864, 26, doi: [10.3847/1538-4357/aad084](https://doi.org/10.3847/1538-4357/aad084)
- Johnston, S., Taylor, R., Bailes, M., et al. 2008, Experimental Astronomy, 22, 151, doi: [10.1007/s10686-008-9124-7](https://doi.org/10.1007/s10686-008-9124-7)
- Jones, E., Oliphant, T. E., Peterson, P., et al. 2001, SciPy: Open source scientific tools for Python. <http://www.scipy.org/>
- Kaspi, V. M., Roberts, M. E., Vasisht, G., et al. 2001, ApJ, 560, 371, doi: [10.1086/322515](https://doi.org/10.1086/322515)
- Kelley, H. J. 1960, ARS Journal, 30, 947, doi: [10.2514/8.5282](https://doi.org/10.2514/8.5282)
- Kiel, P. D., & Hurley, J. R. 2009, MNRAS, 395, 2326, doi: [10.1111/j.1365-2966.2009.14711.x](https://doi.org/10.1111/j.1365-2966.2009.14711.x)
- Kiel, P. D., Hurley, J. R., Bailes, M., & Murray, J. R. 2008, MNRAS, 388, 393, doi: [10.1111/j.1365-2966.2008.13402.x](https://doi.org/10.1111/j.1365-2966.2008.13402.x)
- Kingma, D. P., & Ba, J. 2014, preprint, arXiv:1412.6980. <https://arxiv.org/abs/1412.6980>
- Lai, D., Chernoff, D. F., & Cordes, J. M. 2001, ApJ, 549, 1111, doi: [10.1086/319455](https://doi.org/10.1086/319455)
- Lam, S. K., Pitrou, A., & Seibert, S. 2015, in Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC, LLVM '15 (New York, NY, USA: Association for Computing Machinery). <https://doi.org/10.1145/2833157.2833162>
- Lentati, L., Shannon, R. M., Coles, W. A., et al. 2016, MNRAS, 458, 2161, doi: [10.1093/mnras/stw395](https://doi.org/10.1093/mnras/stw395)
- Levin, L., Bailes, M., Barsdell, B. R., et al. 2013, MNRAS, 434, 1387, doi: [10.1093/mnras/stt1103](https://doi.org/10.1093/mnras/stt1103)
- Li, C., Zhao, G., Jia, Y., et al. 2019, ApJ, 871, 208, doi: [10.3847/1538-4357/aafa17](https://doi.org/10.3847/1538-4357/aafa17)
- Lin, H., Li, X., & Luo, Z. 2020, MNRAS, 493, 1842, doi: [10.1093/mnras/staa218](https://doi.org/10.1093/mnras/staa218)
- Liu, J., Yan, Z., Shen, Z.-Q., et al. 2020, PASJ, 72, 70, doi: [10.1093/pasj/psaa058](https://doi.org/10.1093/pasj/psaa058)
- Lorimer, D. R., & Kramer, M. 2004, Handbook of Pulsar Astronomy (Cambridge University Press)
- Lorimer, D. R., Faulkner, A. J., Lyne, A. G., et al. 2006, MNRAS, 372, 777, doi: [10.1111/j.1365-2966.2006.10887.x](https://doi.org/10.1111/j.1365-2966.2006.10887.x)
- Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, AJ, 129, 1993, doi: [10.1086/428488](https://doi.org/10.1086/428488)
- Manchester, R. N., Lyne, A. G., Camilo, F., et al. 2001, Monthly Notices of the Royal Astronomical Society, 328, 17
- Mandel, I., & Müller, B. 2020, MNRAS, 499, 3214, doi: [10.1093/mnras/staa3043](https://doi.org/10.1093/mnras/staa3043)
- Marchetti, T., Rossi, E. M., & Brown, A. G. A. 2019, MNRAS, 490, 157, doi: [10.1093/mnras/sty2592](https://doi.org/10.1093/mnras/sty2592)
- Matthews, A. M., Nice, D. J., Fonseca, E., et al. 2016, ApJ, 818, 92, doi: [10.3847/0004-637X/818/1/92](https://doi.org/10.3847/0004-637X/818/1/92)
- McKinney, W. 2010, in Proceedings of the 9th Python in Science Conference, ed. S. van der Walt & J. Millman, 51–56. <http://conference.scipy.org/proceedings/scipy2010/mckinney.html>
- Melrose, D. B. 2017, Reviews of Modern Plasma Physics, 1, 5, doi: [10.1007/s41614-017-0007-0](https://doi.org/10.1007/s41614-017-0007-0)
- Miyamoto, M., & Nagai, R. 1975, PASJ, 27, 533
- Morawski, F., & Bejger, M. 2020, A&A, 642, A78, doi: [10.1051/0004-6361/202038130](https://doi.org/10.1051/0004-6361/202038130)
- Motch, C., Pires, A. M., Haberl, F., Schwobe, A., & Zavlin, V. E. 2009, A&A, 497, 423, doi: [10.1051/0004-6361/200811052](https://doi.org/10.1051/0004-6361/200811052)
- Nagakura, H., Sumiyoshi, K., & Yamada, S. 2019, ApJL, 880, L28, doi: [10.3847/2041-8213/ab30ca](https://doi.org/10.3847/2041-8213/ab30ca)
- Narayan, R., & Ostriker, J. P. 1990, ApJ, 352, 222, doi: [10.1086/168529](https://doi.org/10.1086/168529)
- Navarro, J. F., Frenk, C. S., & White, S. D. M. 1996, ApJ, 462, 563, doi: [10.1086/177173](https://doi.org/10.1086/177173)
- Oliphant, T. E. 2006, A guide to NumPy (USA: Trelgol Publishing)
- Ośłowski, S., Bulik, T., Gondek-Rosińska, D., & Belczyński, K. 2011, MNRAS, 413, 461, doi: [10.1111/j.1365-2966.2010.18147.x](https://doi.org/10.1111/j.1365-2966.2010.18147.x)
- Parthasarathy, A., Shannon, R. M., Johnston, S., et al. 2019, MNRAS, 489, 3810, doi: [10.1093/mnras/stz2383](https://doi.org/10.1093/mnras/stz2383)
- Paszke, A., Gross, S., Massa, F., et al. 2019, preprint, arXiv:1912.01703. <https://arxiv.org/abs/1912.01703>
- Pavan, L., Bordas, P., Pühlhofer, G., et al. 2014, A&A, 562, A122, doi: [10.1051/0004-6361/201322588](https://doi.org/10.1051/0004-6361/201322588)
- Perera, B. B. P., DeCesar, M. E., Demorest, P. B., et al. 2019, MNRAS, 490, 4666, doi: [10.1093/mnras/stz2857](https://doi.org/10.1093/mnras/stz2857)

- Perez, F., & Granger, B. E. 2007, *Comput. Sci. Eng.*, 9, 21, doi: [10.1109/MCSE.2007.53](https://doi.org/10.1109/MCSE.2007.53)
- Pichardo, B., Moreno, E., Allen, C., et al. 2012, *AJ*, 143, 73, doi: [10.1088/0004-6256/143/3/73](https://doi.org/10.1088/0004-6256/143/3/73)
- Podsiadlowski, P., Langer, N., Poelarends, A. J. T., et al. 2004, *ApJ*, 612, 1044, doi: [10.1086/421713](https://doi.org/10.1086/421713)
- Polido, P., Jablonski, F., & Lépine, J. R. D. 2013, *ApJ*, 778, 32, doi: [10.1088/0004-637X/778/1/32](https://doi.org/10.1088/0004-637X/778/1/32)
- Rawat, W., & Wang, Z. 2017, *Neural Computation*, 29, 1, doi: [10.1162/NECO_a.00990](https://doi.org/10.1162/NECO_a.00990)
- Reid, M. J., Menten, K. M., Brunthaler, A., et al. 2019, *ApJ*, 885, 131, doi: [10.3847/1538-4357/ab4a11](https://doi.org/10.3847/1538-4357/ab4a11)
- Rozwadowska, K., Vissani, F., & Cappellaro, E. 2021, *NewA*, 83, 101498, doi: [10.1016/j.newast.2020.101498](https://doi.org/10.1016/j.newast.2020.101498)
- Ruder, S. 2017, An overview of gradient descent optimization algorithms. <https://arxiv.org/abs/1609.04747>
- Sartore, N., Ripamonti, E., Treves, A., & Turolla, R. 2010, *A&A*, 510, A23, doi: [10.1051/0004-6361/200912222](https://doi.org/10.1051/0004-6361/200912222)
- Schwab, J., Podsiadlowski, P., & Rappaport, S. 2010, *ApJ*, 719, 722, doi: [10.1088/0004-637X/719/1/722](https://doi.org/10.1088/0004-637X/719/1/722)
- Shklovskii, I. S. 1970, *Soviet Ast.*, 13, 562
- Skliris, V., Norman, M. R. K., & Sutton, P. J. 2020, preprint, arXiv:2009.14611. <https://arxiv.org/abs/2009.14611>
- Skowron, D. M., Skowron, J., Mróz, P., et al. 2019, *Science*, 365, 478, doi: [10.1126/science.aau3181](https://doi.org/10.1126/science.aau3181)
- Smits, R., Kramer, M., Stappers, B., et al. 2009, *A&A*, 493, 1161, doi: [10.1051/0004-6361:200810383](https://doi.org/10.1051/0004-6361:200810383)
- Smits, R., Tingay, S. J., Wex, N., Kramer, M., & Stappers, B. 2011, *A&A*, 528, A108, doi: [10.1051/0004-6361/201016141](https://doi.org/10.1051/0004-6361/201016141)
- Spitzer, Lyman, J. 1942, *ApJ*, 95, 329, doi: [10.1086/144407](https://doi.org/10.1086/144407)
- Stovall, K., Lynch, R. S., Ransom, S. M., et al. 2014, *ApJ*, 791, 67, doi: [10.1088/0004-637X/791/1/67](https://doi.org/10.1088/0004-637X/791/1/67)
- Tamborra, I., Hanke, F., Janka, H.-T., et al. 2014, *ApJ*, 792, 96, doi: [10.1088/0004-637X/792/2/96](https://doi.org/10.1088/0004-637X/792/2/96)
- Tauris, T. M., Langer, N., Moriya, T. J., et al. 2013, *ApJL*, 778, L23, doi: [10.1088/2041-8205/778/2/L23](https://doi.org/10.1088/2041-8205/778/2/L23)
- Tauris, T. M., & Manchester, R. N. 1998, *MNRAS*, 298, 625, doi: [10.1046/j.1365-8711.1998.01369.x](https://doi.org/10.1046/j.1365-8711.1998.01369.x)
- Tauris, T. M., Kaspi, V. M., Breton, R. P., et al. 2015, in *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, 39. <https://arxiv.org/abs/1501.00005>
- Taylor, J. H., & Cordes, J. M. 1993, *ApJ*, 411, 674, doi: [10.1086/172870](https://doi.org/10.1086/172870)
- Vallée, J. P. 2017, *Astron. Rev.*, 13, 113, doi: [10.1080/21672857.2017.1379459](https://doi.org/10.1080/21672857.2017.1379459)
- van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *Comput. Sci. Eng.*, 13, 22, doi: [10.1109/MCSE.2011.37](https://doi.org/10.1109/MCSE.2011.37)
- Verbunt, F., Igoshev, A., & Cator, E. 2017, *A&A*, 608, A57, doi: [10.1051/0004-6361/201731518](https://doi.org/10.1051/0004-6361/201731518)
- Viganò, D., Rea, N., Pons, J. A., et al. 2013, *MNRAS*, 434, 123, doi: [10.1093/mnras/stt1008](https://doi.org/10.1093/mnras/stt1008)
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nature Methods*, 17, 261, doi: [10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2)
- Wainscoat, R. J., Cohen, M., Volk, K., Walker, H. J., & Schwartz, D. E. 1992, *ApJS*, 83, 111, doi: [10.1086/191733](https://doi.org/10.1086/191733)
- Walter, F. M., Eisenbeiß, T., Lattimer, J. M., et al. 2010, *ApJ*, 724, 669, doi: [10.1088/0004-637X/724/1/669](https://doi.org/10.1088/0004-637X/724/1/669)
- Wang, J. B., Coles, W. A., Hobbs, G., et al. 2017, *MNRAS*, 469, 425, doi: [10.1093/mnras/stx837](https://doi.org/10.1093/mnras/stx837)
- Wei, W., & Huerta, E. A. 2020, preprint, arXiv:2010.09751. <https://arxiv.org/abs/2010.09751>
- Wong, K. W. K., & Gerosa, D. 2019, *PhRvD*, 100, 083015, doi: [10.1103/PhysRevD.100.083015](https://doi.org/10.1103/PhysRevD.100.083015)
- Yadan, O. 2019, *Hydra - A framework for elegantly configuring complex applications*, Github. <https://github.com/facebookresearch/hydra>
- Yao, J. M., Manchester, R. N., & Wang, N. 2017, *ApJ*, 835, 29, doi: [10.3847/1538-4357/835/1/29](https://doi.org/10.3847/1538-4357/835/1/29)
- Yusifov, I., & Küçük, I. 2004, *A&A*, 422, 545, doi: [10.1051/0004-6361:20040152](https://doi.org/10.1051/0004-6361:20040152)

Table 5 (continued)

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.
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	1.681114 × 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-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	08:24:24.8022(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 ± 0.1	-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	09:00:43.93088(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(6)	1.912	PSR	
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)	0.862	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	10.2521(1)	10.2521(1)	0.834	PSR	
J1023+0038	10:23:47.08720(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	
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	4.4663 × 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	10:45:00.18696(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	0.08990 × 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	3.738387 × 10 ⁻¹⁵	4.84066(34)	0.414	PSR	
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	43.4233(7)	1.244	Binary PSR	
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	-	49620	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	

Table 5 continued

Table 5 (continued)

JName	RA [hms]	DEC [$^{\circ}$, ' , '']	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.
J2219+2933	22:19:23.3444(1)	+29:33:05.411(1)	-5.53 \pm 0.08	-11.34 \pm 0.11	0.23 \pm 0.09	49686	1.00459	4.9517 \times 10 $^{-16}$	74.5213(15)	25	PSR	
J2214+3000	22:14:38.8525(57)	+30:00:38.193(97)	20.7 \pm 0.1	-1.6 \pm 0.2	1.1 \pm 0.2	56610	0.00312	1.47261 \times 10 $^{-20}$	22.5450(25)	1.673	PSR	[6]
J2215+5135	22:15:32.6864(1)	+51:35:36.4063(1)	0.31 \pm 0.54	1.88 \pm 0.6	0.28 \pm 0.36	57205.875	0.00261	3.34 \times 10 $^{-20}$	69.1951(2)	2.773	Binary PSR	[5]
J2219+4754	22:19:48.139(3)	+47:54:53.93(3)	-12 \pm 8	-30 \pm 6	-	46599	0.53847	2.765209 \times 10 $^{-15}$	43.4975(6)	2.387	PSR	
J2222+0137	22:22:05.96910(17)	-01:37:15.72441(10)	44.73 \pm 0.02	-5.68 \pm 0.06	3.742 \pm 0.01	55473	0.03282	5.8024 \times 10 $^{-20}$	3.277	0.267	PSR	
J2225+6335	22:25:52.8627(3)	+65:35:36.371(1)	147.23 \pm 0.23	126.5 \pm 0.1	1.2 \pm 0.2	53880	0.68254	9.6606 \times 10 $^{-15}$	36.44362(51)	1.881	PSR	
J2229+2643	22:29:50.8845(1)	+26:43:57.648(13)	-1.7 \pm 0.1	-5.9 \pm 0.1	0.8 \pm 0.6	56937	0.00298	1.516 \times 10 $^{-21}$	22.7271(6)	1.8	PSR	[6]
J2234+0611	22:34:23.073090(2)	+06:11:28.68633(7)	25.3 \pm 0.02	9.71 \pm 0.05	1.03 \pm 0.04	56794.09	0.00358	1.201419 \times 10 $^{-20}$	10.7645(15)	0.854	PSR	
J2235+1506	22:35:43.7001(1)	+15:06:49.060(5)	15 \pm 4	10 \pm 8	-	49250	0.05977	1.576 \times 10 $^{-19}$	18.09(5)	1.54	PSR	
J2240+5832	22:40:42.939(4)	+58:32:37.52(4)	-6.1 \pm 0.8	-21 \pm 0.4	-	54000	0.13993	1.5358 \times 10 $^{-14}$	263.50(5)	7.275	PSR	
J2241+5236	22:41:42.01850(5)	-52:36:36.2260(3)	17.1 \pm 0.1	-3.32 \pm 0.05	-	55044.2	0.00219	6.64 \times 10 $^{-21}$	11.41151(2)	0.963	PSR	
J2248-0101	22:48:26.8859(1)	-01:01:48.085(1)	-10.5 \pm 0.07	-17.49 \pm 0.19	0.25 \pm 0.06	50866	0.47723	6.5935 \times 10 $^{-16}$	29.05(3)	25	PSR	
J2301+5852	23:01:08.29(33)	+58:52:44.45(60)	-6.4 \pm 0.6	-2.3 \pm 0.6	-	54852	6.97907	4.7123 \times 10 $^{-13}$	-	3.3*	MAG	
J2302+4442	23:02:46.9783(78)	+44:42:22.085(95)	-5 \pm 4	-7 \pm 5	1 \pm 1	56675	0.00519	1.378 \times 10 $^{-20}$	13.7279(3)	0.859	PSR	[6]
J2305+3100	23:05:58.3212(1)	+31:00:01.281(1)	-3.7 \pm 0.04	-15.63 \pm 0.11	0.23 \pm 0.03	48714	1.57589	2.89250 \times 10 $^{-15}$	49.5845(12)	4.35*	PSR	
J2308+5547	23:08:13.822(5)	+55:47:36.03(4)	-15 \pm 8	0 \pm 27	-	48717	0.47507	1.99491 \times 10 $^{-16}$	46.53905(37)	2.075	PSR	
J2313+4253	23:13:08.6209(1)	+42:53:13.043(2)	24.15 \pm 0.1	5.95 \pm 0.13	0.93 \pm 0.07	52275	0.34943	1.123699 \times 10 $^{-16}$	17.27693(33)	1.108	PSR	
J2317+1439	23:17:09.2364(1)	+14:39:31.265(1)	-1.33 \pm 0.3	3.51 \pm 0.5	0.1 \pm 0.1	54977	0.00345	2.43012 \times 10 $^{-21}$	21.877(9)	2.16	PSR	[6]
J2317+2149	23:17:57.8419(1)	+21:49:48.019(1)	8.49 \pm 0.07	0.19 \pm 0.14	0.51 \pm 0.05	48716	1.44465	1.047047 \times 10 $^{-15}$	20.86959(31)	1.799	PSR	
J2321+6024	23:21:55.20(73)	+60:24:31.0(55)	-17 \pm 22	-7 \pm 19	-	49303	2.25649	7.03689 \times 10 $^{-15}$	94.591(18)	2.513	PSR	
J2322+2057	23:22:22.33516(7)	+20:57:02.6772(14)	-18.5 \pm 0.3	-15.3 \pm 0.5	-	55000	0.00481	9.661 \times 10 $^{-20}$	13.36(4)	1.009	PSR	[6]
J2322+2650	23:22:34.64004(3)	-26:50:58.3171(6)	-2.4 \pm 0.2	-8.3 \pm 0.4	4.4 \pm 1.2	56152	0.00346	5.834 \times 10 $^{-22}$	6.149(2)	0.76	PSR	
J2325+6316	23:25:13.3196(2)	+63:16:52.362(1)	-5.92 \pm 0.08	-2.05 \pm 0.19	0.01 \pm 0.05	48309	1.43631	2.82546 \times 10 $^{-15}$	197.37(5)	4.855	PSR	
J2330+2005	23:30:26.885(7)	-20:05:29.63(17)	74.7 \pm 1.9	5 \pm 3	-	49878	1.64362	4.63038 \times 10 $^{-15}$	8.456(2)	0.861	PSR	
J2337+6151	23:37:05.762(1)	+61:51:01.53(7)	-1 \pm 18	-15 \pm 16	-	53100	0.49537	1.934498 \times 10 $^{-13}$	58.410(15)	2.079	PSR	
J2339+0533	23:39:38.74(1)	-05:33:05.32(3)	4.15 \pm 0.48	-10.31 \pm 0.3	0.75 \pm 0.25	55100	0.00288	1.4102 \times 10 $^{-20}$	8.717(4)	0.751	Binary PSR	
J2346-0609	23:46:50.4978(1)	-06:09:59.899(2)	37.38 \pm 0.03	-20.21 \pm 0.09	0.27 \pm 0.03	51021	1.18146	1.36305 \times 10 $^{-15}$	22.504(19)	3.7*	PSR	
J2354+6155	23:54:04.7830(1)	+61:55:46.845(1)	22.76 \pm 0.05	4.897 \pm 0.025	0.41 \pm 0.04	49405	0.94478	1.626279 \times 10 $^{-14}$	94.662(6)	2.399	PSR	