

Cataloging accreted stars within *Gaia* DR2 using deep learning

B. Ostdiek^{1,*}, L. Necib², T. Cohen¹, M. Freytsis^{3,4}, M. Lisanti⁵, S. Garrison-Kimmelman⁶, A. Wetzel⁷, R. E. Sanderson^{8,9}, and P. F. Hopkins⁶

¹ Institute of Theoretical Science, Department of Physics, University of Oregon, Eugene, OR 97403, USA
e-mail: bostdiek@g.harvard.edu

² Walter Burke Institute for Theoretical Physics, California Institute of Technology, Pasadena, CA 91125, USA
e-mail: lnebib@caltech.edu

³ Raymond and Beverly Sackler School of Physics and Astronomy, Tel Aviv University, Tel Aviv 69978, Israel

⁴ School of Natural Sciences, Institute for Advanced Study, Princeton, NJ 08540, USA

⁵ Department of Physics, Princeton University, Princeton, NJ 08544, USA

⁶ TAPIR, California Institute of Technology, Pasadena, CA 91125, USA

⁷ Department of Physics, University of California, Davis, CA 95616, USA

⁸ Department of Physics and Astronomy, University of Pennsylvania, Philadelphia, PA 19104, USA

⁹ Center for Computational Astrophysics, Flatiron Institute, New York, NY 10010, USA

Received 7 October 2019 / Accepted 21 January 2020

ABSTRACT

Aims. The goal of this study is to present the development of a machine learning based approach that utilizes phase space alone to separate the *Gaia* DR2 stars into two categories: those accreted onto the Milky Way from those that are in situ. Traditional selection methods that have been used to identify accreted stars typically rely on full 3D velocity, metallicity information, or both, which significantly reduces the number of classifiable stars. The approach advocated here is applicable to a much larger portion of *Gaia* DR2.

Methods. A method known as “transfer learning” is shown to be effective through extensive testing on a set of mock *Gaia* catalogs that are based on the FIRE cosmological zoom-in hydrodynamic simulations of Milky Way-mass galaxies. The machine is first trained on simulated data using only 5D kinematics as inputs and is then further trained on a cross-matched *Gaia*/RAVE data set, which improves sensitivity to properties of the real Milky Way.

Results. The result is a catalog that identifies $\sim 767\,000$ accreted stars within *Gaia* DR2. This catalog can yield empirical insights into the merger history of the Milky Way and could be used to infer properties of the dark matter distribution.

Key words. Galaxy: kinematics and dynamics – Galaxy: halo – solar neighborhood – catalogs – methods: data analysis

1. Introduction

The two dominant theories of Galactic formation have been monolithic collapse (Eggen et al. 1962) and a slower build up of merging protogalaxies (Searle & Zinn 1978). The compatibility of observations with the Lambda–cold dark matter (Λ CDM) model, augmented by cosmological simulations, provides strong support for the hierarchical structure formation hypothesis (White & Rees 1978). The now prevailing view is that a slow and complex process of accretion provides the majority of the dark matter in a galaxy and builds up the stellar halo (Helmi & White 1999; Bullock et al. 2001; Bullock & Johnston 2005). There is abundant evidence for such accretion events. For example, the Sagittarius dwarf spheroidal galaxy is currently being disrupted into a large stream that fills the sky (Ibata et al. 1994; Majewski et al. 2003). The Large and Small Magellanic Clouds (Weinberg 1998; Jiang & Binney 1999), as well as the Field of Streams (Belokurov et al. 2006, 2007), are easily accommodated within this accretion driven framework.

The second release of data from the *Gaia* satellite (DR2) (Gaia Collaboration et al. 2018a) allows for the identification of more imprints of accretions, further refining the picture. It provides parallax and proper motion measurements for over 1.3

billion stars, with high-quality 6D phase-space measurements available for a local subset of over 5 million stars, these multi-dimensional measurements open the door to a deeper understanding of the Milky Way.

Studies of the *Gaia* data have already revealed new structures identified as the remnants of the complicated Milky Way accretion history, including *Gaia* Enceladus (a.k.a. the *Gaia* Sausage) (Belokurov et al. 2018; Helmi et al. 2018; Myeong et al. 2018a; Lancaster et al. 2019) and Sequoia (Myeong et al. 2019), as well as an abundance of new streams (Koppelman et al. 2018; Myeong et al. 2018b; Malhan & Ibata 2018) and disrupted dwarf galaxies (Fernández-Trincado et al. 2019a,b). Unsurprisingly, searches for such accretion events can be highly non-trivial, as they require identifying clusters of accreted stars that share the same origin against the background of the stellar disk.

There have been recent attempts to automate finding structure in the *Gaia* data¹. For instance, assuming a gravitational potential, Malhan & Ibata (2018) find members of streams contained in 6D flux tubes. Model independent searches have been done with clustering on the integrals of motion (Borsato et al. 2019) which also needs the full 6D phase space. Both of these tech-

* Current address: Department of Physics, Harvard University, Cambridge, MA 02138, USA.

¹ See, e.g., Mateu et al. (2011) for related earlier techniques that combine regression techniques with physically-motivated data preprocessing.

niques need to have their halo and accreted stars already identified. [Veljanoski et al. \(2019\)](#) use a boosted decision tree to classify halo stars, which needs to have metallicity and the full phase space information to do well. This motivates the use of so-called deep learning techniques, as they are capable of finding non-linear relationships in data using unprocessed low-level inputs (see e.g., [Carleo et al. 2019](#), for a recent review of deep-learning applications in the physical sciences). Applications to the Large Hadron Collider (e.g., [Larkoski et al. 2020](#)), the Large Synoptic Survey Telescope (e.g., [Hezaveh et al. 2017](#)), many-body quantum systems (e.g., [Zhang et al. 2019](#)), and galaxy morphology (e.g., [Huertas-Company et al. 2018; Domínguez Sánchez et al. 2019](#)) make a clear case for the usefulness of deep learning across many disciplines. The power of these algorithms is most evident when used on large datasets that contain many non-trivial correlations that can be leveraged to expose interesting structure.

In this paper, we use deep learning to identify accreted stars within *Gaia* DR2, which will allow us to boost the size of our sample. We build a deep neural-network-based classifier, trained on a carefully curated combination of state-of-the-art cosmological simulations augmented by a subset of the *Gaia* data itself. When applied to *Gaia* DR2, the result is a high-purity catalog of accreted Milky Way stars, which can be used to search for new structures within our Galaxy.

Of crucial importance to the classifiers developed here are the mock catalogs of [Sanderson et al. \(2020\)](#). Because it is possible to identify the accreted population at truth level within these simulated datasets, they allow us to test and optimize our approach before applying the methods to real data. Additionally, we rely on these mock catalogs to provide the first stage of training, preparing the network for a further stage of training that utilizes actual Milky Way data. Many efforts in building such mock catalogs already exist. For example, Galaxia ([Sharma et al. 2011](#)) samples mock stellar halos from [Bullock & Johnston \(2005\)](#) and mock catalogs based on resampling stars from cosmological simulations have been developed in [Bahcall & Soneira \(1980\)](#), [Robin & Creze \(1986\)](#), [Bienayme et al. \(1987\)](#), [Lowing et al. \(2015\)](#), [Grand et al. \(2018\)](#). However, with the increase in resolution of hydrodynamic simulations ([Marinacci et al. 2014; Wang et al. 2015; Fattahi et al. 2016; Wetzel et al. 2016; Grand et al. 2017; Kelley et al. 2019](#)) and the improved modeling of baryonic physics, simulations of the Milky Way can help make self-consistent predictions for survey outputs, especially by utilizing the full knowledge of the formation history.

In this work, we focus on the details of how to use deep neural networks to build a catalog of accreted stars in *Gaia*. The first science results using this catalog are presented in [Necib et al. \(2019a\)](#), which uses our catalog to reproduce known structures including *Gaia* Enceladus and the Helmi stream and identify potential undiscovered structures. In particular, we identify a previously-unknown massive stream, Nyx, in the vicinity of the Sun comprising more than 10% of local accreted stars. A detailed study of its known properties is presented in [Necib et al. \(2019b\)](#).

This paper is detailed and mostly self-contained. As such, some readers may prefer to pass over certain sections. Section 2 briefly reviews the FIRE simulations, with an emphasis on the *Gaia* mock catalogs. Our working definition of accreted stars is given in Sect. 2.2. Section 3 reviews the traditional selection methods utilized for identifying accreted stars, which will be useful benchmarks to compare to the approach developed here. Section 4 provides the machine learning architecture and discusses the training procedure as well as the incorporation of measurement errors. Section 5 studies what observables should be used by the neural network. Section 5.1 finds the opti-

mal observables, assuming all stars are measured well. Next, Sect. 5.2 shows how this changes when trying to generalize to dimmer stars. Section 5.3 demonstrates that the network can generalize to new view points (as a proxy for new galaxies). Section 6 briefly explores what physical characteristics of the data the machine is using to identify accreted stars. Section 7 introduces the idea of transfer learning, which eventually allows us to train on stars in the Milky Way. The overall strategy is discussed in Sect. 7.1. In Sect. 7.2, we find the optimal empirical labels to use for the transfer process. In addition, we present what network score cut leads to the most consistent results. The strategy and cut are then validated on an independent FIRE catalog in Sect. 7.3. Section 8 describes the application to *Gaia* data, producing the catalog of accreted stars. We first validate the procedure on RAVE DR5-*Gaia* DR2 cross-matched stars in Sect. 8.1. The final accreted catalog is presented in Sect. 8.2. Section 9 contains the conclusion. Appendix A shows the stars which the network classifies as in situ. Appendix B provides validation of the error sampling procedure used during the training of the neural network. Appendix C provides an alternative catalog by allowing the network to use photometry as well as kinematics, which is shown in Sect. 5.2 not to generalize as well. We use the Matplotlib ([Hunter 2007](#)), Numpy ([van der Walt et al. 2011](#)), pandas ([McKinney 2010](#)), scikit-learn ([Pedregosa et al. 2011](#)), and scipy ([Jones et al. 2001](#)) throughout our analysis.

2. Simulating the *Gaia* Sky

Our goal is to develop methods by which we can obtain a catalog of accreted stars. To do so, we build and test our methodology using mock *Gaia* catalogs derived from the *Latte* suite of FIRE-2 cosmological hydrodynamic simulations of Milky Way-mass halos ([Wetzel et al. 2016; Hopkins et al. 2018a](#)) – specifically, those obtained from the simulated galaxies named m12i and m12f from [Sanderson et al. \(2020\)](#). This allows us to train and validate our approach on samples where every star's true history is known ([Necib et al. 2019c](#)). Below, we briefly review the physics underlying the mocks and describe the algorithm used to identify accreted stars.

2.1. FIRE Simulations and the *Gaia* Mock Catalogs

The numerical methods and physics in the simulations are presented in extensive detail in [Hopkins et al. \(2018a\)](#). The FIRE-2 simulations are run with GIZMO², a multi-method gravity and hydrodynamics code ([Hopkins 2015](#)), using a hybrid tree-PM gravity solver (see e.g., [Springel 2005](#)) and the meshless finite-mass (MFM) Lagrangian Godunov solver for hydrodynamics. Radiative heating/cooling is included over the temperature range $10\text{--}10^{10}$ K, assuming a uniform but redshift-dependent metagalactic UV background from [Faucher-Giguère et al. \(2009\)](#), as well as metal-line, molecular, and other processes. Star formation occurs only in self-gravitating ([Hopkins et al. 2013](#)), Jeans unstable, self-shielding, molecular ([Krumholz & Gnedin 2011](#)) gas at densities over 1000 cm^{-3} . Once stars form, all feedback rates are calculated using standard stellar evolution models (e.g., [Leitherer et al. 1999](#)), assuming a [Kroupa \(2001\)](#) initial mass function and the known age, metallicity, and mass of the star particle. The simulations explicitly model mechanical feedback from supernovæ (Ia and core-collapse) and stellar mass loss (O/B and AGB) ([Hopkins et al. 2018b](#)) as well as multi-wavelength radiative feedback including photoionization

² A public version of GIZMO can be found at <http://www.tapir.caltech.edu/~phopkins/Site/GIZMO.html>.

and photoelectric heating and radiation pressure from single and multiple scattering (Hopkins et al. 2020).

Our study relies on two Milky Way-mass galaxies from the *Latte* suite of FIRE-2 simulations, m12i and m12f, introduced in Wetzel et al. (2016), Garrison-Kimmel et al. (2017). The simulations adopt a flat Λ CDM cosmology with $(\Omega_m, \Omega_b, h, \sigma_8, n_s) = (0.272, 0.0455, 0.702, 0.807, 0.961)$, with a high-resolution “zoom-in” region (e.g., Oñorbe et al. 2014), a fully uncontaminated diameter greater than 1.2 Mpc at $z = 0$, and a baryonic mass resolution $\approx 7000 M_\odot$ surrounding the halo of interest in a large cosmological box. The galaxies m12i and m12f have present-day stellar masses of 5.5 and $6.9 \times 10^{10} M_\odot$, respectively, comparable to the Milky Way mass of $(5 \pm 1) \times 10^{10} M_\odot$ (Bland-Hawthorn & Gerhard 2016). More detailed stellar halo structure comparisons to the Milky Way for both have been previously presented in Bonaca et al. (2017), Sanderson et al. (2018).

Mock catalog generation is described in Sanderson et al. (2020). Each star particle of mass $\approx 7000 M_\odot$ is treated as a single-age, single-metallicity population: individual stars are generated by populating the initial mass function according to the particle properties and distributing them within the star particle volume in phase space. To preserve the wide dynamic range of phase-space densities, the phase-space smoothing kernel is subdivided into 8 age bins for stars formed in situ, while a separate kernel is used for all accreted stars. Three different solar positions, each defining a different local standard of rest (LSR), are chosen to construct mocks, spaced uniformly around a circle $R_\odot = 8.2$ kpc from the galaxy center, defined by the total angular momentum of disk stars. Lines-of-sight to each star are ray-traced from the solar positions, including extinction and reddening computed self-consistently from the dust and gas in the simulation (and convolved with realistic measurement errors) to compute the observed photometry.

The simulations are not tuned to the Milky Way other than by choosing a dark matter halo of roughly the same mass and choosing from a larger suite of simulations those with similar disk-to-bulge ratios (Garrison-Kimmel et al. 2018) and total stellar masses. They are, therefore, not perfect Milky Way analogues. While the total mass (see above) and radial distribution of stars in the disk are quite similar (e.g., the radius enclosing 90% of the stellar mass within $|z| < 1.1$ kpc is 2.7 kpc (3.4 kpc) in m12i (m12f), compared to 2.6 ± 0.5 kpc in the Milky Way), the stellar velocity dispersion and scale-height at R_\odot are somewhat larger (e.g., the volume-density of stars within $|z| < 200$ pc for both is $\sim 20 M_\odot \text{ pc}^{-3}$, versus $\sim 40 M_\odot \text{ pc}^{-3}$ for the Milky Way from Bland-Hawthorn & Gerhard 2016). Likewise, present-day galaxy-wide star formation rates in simulated galaxies are $\sim 3\text{--}8 M_\odot \text{ yr}^{-1}$, compared to $\sim 1.5\text{--}3 M_\odot \text{ yr}^{-1}$ in the Milky Way (Chomiuk & Povich 2011).

2.2. Identifying accreted simulated stars

We derive truth labels for the simulated stars using the method of Necib et al. (2019c). Our algorithm for tracking the origin of these stars begins by first locating the ones within a Galactocentric distance $r_{\text{GC}} < 16$ kpc at the present day. Stars that do not pass this cut have poor kinematic resolution in *Gaia* DR2, and are therefore not of interest here. For our purposes, accreted stars are defined as bound to a dark matter subhalo that fell into the galaxy. To find such stars, we first define subhalos by applying the ROCKSTAR phase-space finder³ (Behroozi et al. 2013) to the dark matter particles alone. We assign stars to these subhalos by then working

backwards from the present day to redshift $z = 99$, spanning 600 individual snapshots of the simulation. At each snapshot, all stars within the virial radius R_{200m} whose velocity is within 2.5σ of the velocity dispersion for a given subhalo are associated with that subhalo. If a star is associated with a particular subhalo (other than the main halo) for three out of four consecutive snapshots, it is designated as belonging to that subhalo.

We label these star particles as accreted, and label those that do not pass this selection as in situ. To validate this procedure, we also use an alternative approach and apply a cut on the formation distance from the central galaxy at 20 kpc (25 kpc) on m12i (m12f) (see Sanderson et al. 2018, Fig. 1). This results in a similar categorization into accreted and in situ stars. Specifically, for the stars identified by either of these methods, $\sim 74\%$ (65%) are selected by both, $\sim 7\%$ (17%) are selected only by the high formation distance, and the remaining $\sim 19\%$ (18%) are selected only using the merger history for m12i (m12f)⁴. In the rest of the study, we use the merger history to define whether or not a star has been accreted.

3. Traditional selection criteria

The primary task of this work is to develop a neural network to distinguish between accreted and in situ stars. To benchmark the success of such networks, we will compare them to standard selection methods used in the literature. The “traditional” approaches take advantage of the fact that the origin of a star is known to be correlated with its stellar position, velocity, and chemical composition. Relying on simple models for stellar distributions with respect to these inputs, one can motivate a set of selection criteria to identify accreted stars. In what follows, we benchmark the performance of the machine against three specific traditional approaches, which we refer to as the V (e.g., Nissen & Schuster 2010), VM (e.g., Helmi et al. 2017; Posti et al. 2018) and ZM selections (e.g., Herzog-Arbeitman et al. 2018; Necib et al. 2019c). The methods, along with their names, are explained in more detail below. Coordinate transformations are performed with Astropy (Astropy Collaboration 2018).

In Fig. 1, we plot Toomre diagrams of the truth-level velocity distributions of accreted and in situ stars for LSR0 of m12i before applying measurement uncertainties. These plots present the three components of velocity in a Cartesian coordinate system where v_x points from the Galactic center to the sun, v_y points along the direction of the rotation of the Galactic disk, and v_z points towards the angular momentum vector of the disk. We require that all stars in the mock catalog have a small parallax error ($\delta\varpi/\varpi \leq 0.10$) and a measured radial velocity. We note that for the mock datasets, the parallax error cut also acts as an effective distance cut and removes stars farther than ~ 4.5 kpc away. The in situ population is expected to rotate with the disk, consistent with its highest density being near $v_{\text{LSR}} = (0, 224, 0) \text{ km s}^{-1}$. The accreted population, on the other hand, is relatively uniformly distributed in the Toomre diagram. It is, however, worth noting the presence of structure in the accreted population, which could be the imprint of particular merger events.

In what follows, we frequently refer to the “purity” and accreted “efficiency” of a given selection criterion. By purity we mean the fraction of stars defined by a given selection criterion

⁴ The merger history criterion is likely more inclusive due to some stars forming in satellites orbiting within 20–25 kpc. The distance cut would identify these as in situ, even though they formed in a different subhalo.

³ <https://bitbucket.org/pbehroozi/rockstar-galaxies>

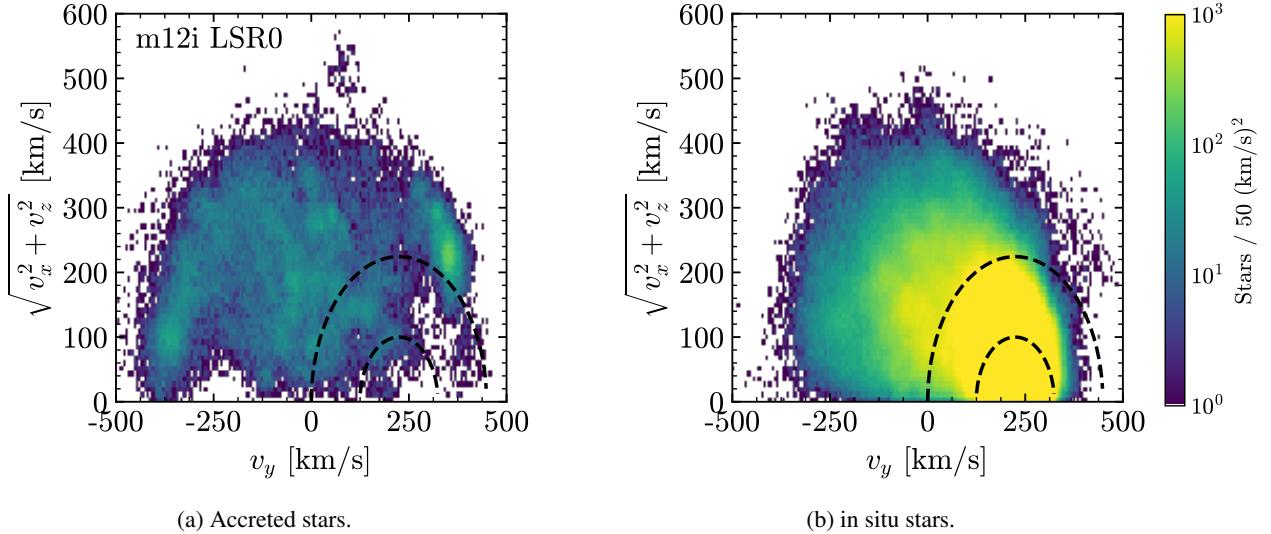


Fig. 1. Truth-level Toomre distributions for LSRO of m12i with $\delta\varpi/\varpi < 0.10$ and measured v_r . Bright yellow regions denote higher density while the darker blue regions denote lower density, with density ranging from 1– 10^3 stars/50 (km s^{-1}) 2 . The outer (inner) dashed lines denote $|\mathbf{v} - \mathbf{v}_{\text{LSR}}| > 224 \text{ km s}^{-1}$ (124 km s^{-1}) with $\mathbf{v}_{\text{LSR}} = (0, 224, 0) \text{ km s}^{-1}$. Velocities are given in Cartesian coordinates centered on the sun, with the x axis oriented away from the Galactic center. Accreted stars are essentially uniformly distributed in this plane, although with structure clearly visible, while the in situ stars peak towards \mathbf{v}_{LSR} , as expected.

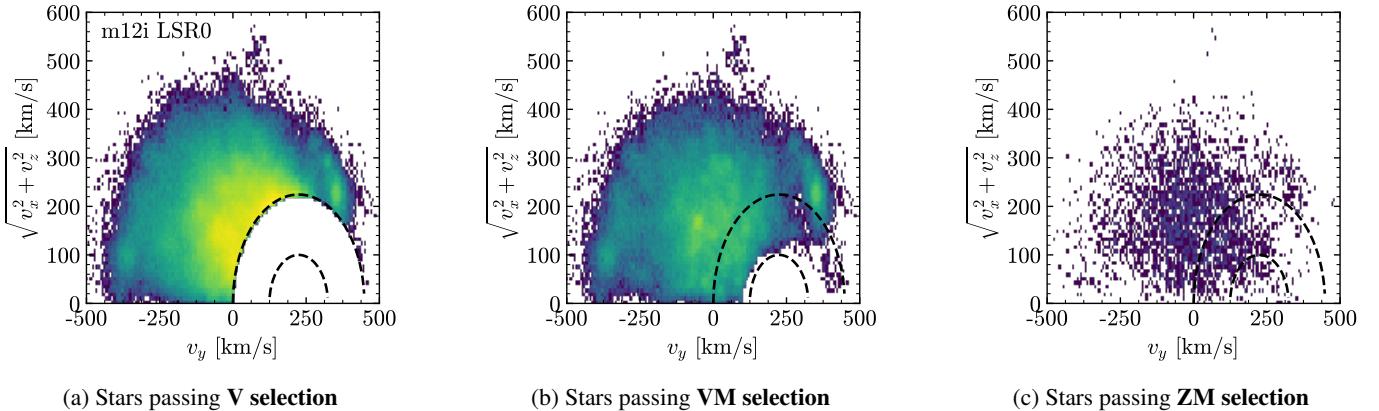


Fig. 2. Toomre diagrams for m12i LSRO for stars passing traditional cut-based selection criteria, see Fig. 1 for the truth-level plots. The velocity based criterion in Figs. 1a and 1b introduces an asymmetry in the Toomre plane, which is absent in Fig. 1c. Only the V selection is possible with *Gaia* DR2 without cross-matching with other catalogs. The purities and efficiencies of all methods are presented in Table 1.

that carry an accreted truth label:

$$\text{Purity} = \frac{N_{\text{accreted, selected}}}{N_{\text{selected}}}. \quad (1)$$

Efficiency is defined as the fraction of all available truth-level stars that are identified, with accreted (ϵ_A) and in situ efficiency (ϵ_I) defined as:

$$\epsilon_A = \frac{N_{\text{accreted, selected}}}{N_{\text{accreted}}}, \quad (2)$$

$$\epsilon_I = \frac{N_{\text{insitu, selected}}}{N_{\text{insitu}}}.$$

A larger ϵ_A means that a greater fraction of accreted stars are selected, but one does not want to do this at the expense of picking in situ stars. The purity is a measure of this contamination.

A simple approach to separating the two stellar populations is to apply the V selection criterion (Nissen & Schuster 2010): a star with velocity \mathbf{v} is defined as accreted if $|\mathbf{v} - \mathbf{v}_{\text{LSR}}| >$

$|\mathbf{v}_{\text{LSR}}| = 224 \text{ km s}^{-1}$. Figure 2a shows the accreted distribution inferred using this criterion. This method clearly biases the selected accreted population in velocity space as it aggressively removes all stars with velocities close to \mathbf{v}_{LSR} , while still presenting significant contamination near the cutoff. As a result, in this case we find a purity of 9.6%, but with an accreted efficiency of $\epsilon_A = 86\%$ due to the relatively inclusive selection criterion. As this method utilizes the 3D velocity, it is only applicable for stars that have line-of-sight velocity measurements in addition to the proper motion measurements.

To improve upon this kinematic-only selection criterion, one can use stellar metallicity as an additional handle. This is the reasoning underlying the method we refer to as the velocity and metallicity or VM selection (Helmi et al. 2017; Posti et al. 2018). This approach is implemented in two stages. First, the 3D velocity distribution of the stars is fit to a two-component Gaussian mixture model. One group has a peak consistent with the disk (again \mathbf{v}_{LSR}). The accreted population is defined to be stars having $[\text{Fe}/\text{H}] < -1$ and more likely to belong to the Gaussian component whose peak velocity is not consistent with

Table 1. Performance metrics for the traditional selection methods.

Method	Purity	ϵ_A	ϵ_I	v_{los}	[Fe/H]
V selection	9.6%	86.0%	4.5%	✓	
VM selection	30.5%	86.2%	1.1%	✓	✓
ZM selection	50.9%	2.4%	0.02%	✓	

Notes. The middle columns come from the truth-level information of the m12i LSR0 catalog with $\delta\varpi/\varpi \leq 0.10$. The V and VM selections require a measurement of the radial velocity, denoted as v_{los} . The purity of a selection is the fraction of selected stars with accreted truth labels. The accreted efficiency (ϵ_A) is the fraction of accreted stars in the sample that get selected and the in situ efficiency (ϵ_I) is the fraction of in situ stars selected, making ϵ_I^{-1} the background rejection factor. The right columns show the measurements primarily responsible for limiting available statistics, but which are required to use the method.

the disk. As is evident from Fig. 2b, this selection criteria allows one to identify accreted stars whose kinematics are more similar to those of the disk. The accreted efficiency of the resulting selection is again $\epsilon_A = 86\%$ but now with a purity of 30.5%. Although the fraction of falsely labeled accreted stars decreases as compared to the V selection, they still make up a majority of the selected sample. Relative to the V selection, this method can only be applied to a smaller fraction of stars as it requires a metallicity measurement; see Table 8 which provides the relevant statistics for *Gaia* DR2.

Both of these approaches demonstrate the intuitive fact that hard cuts on stellar velocities will place a bias on the kinematics of the resulting sample. Therefore, one must take great care when using the resulting samples as the input to any study that aims to reconstruct the velocity distribution of the accreted stellar population. An alternative approach is to separate out the in situ population based on the position and metallicity of the stars alone. This motivates the z and metallicity or ZM selection (Herzog-Arbeitman et al. 2018; Necib et al. 2019c), which identifies accreted stars by selecting for those with $|z| > 1.5$ kpc and $[\text{Fe}/\text{H}] < -1.5$. As is clear by inspecting Fig. 2c, this requirement is the most conservative, with an accreted efficiency of only 2.4%. However, for the same mock *Gaia* catalog, it leads to a purity of 50.9% with no obvious bias visible in the Toomre diagram.

The performance of these methods, summarized in Table 1, is a testament to the challenge of trying to develop unbiased accreted star selection criteria while maintaining high purity and efficiency. A price is paid in statistics to apply the kinematic selection criteria of these methods. Specifically, in *Gaia* DR2, only 0.3% of stars have accurate 6D phase-space measurements. Cross-matching with other observations is required to get metallicities, reducing this to merely 0.02% of the full dataset in the case of RAVE DR5, see Table 8 for context. However, it is worth emphasizing that these issues persist even if full 3D velocity and metallicity measurements are available. One of the primary benefits of our methodology is its relative flexibility with respect to the completeness of stellar data available. As we demonstrate below, the machine learning algorithms employed here achieve comparable or better purity of accreted stars, even in cases where the full 6D phase space and/or chemical composition is not available.

4. Machine learning on FIRE

Our goal is to identify accreted stars even in cases where the full phase-space information and/or chemical abundances are

unavailable. To this end, we train a number of dense multi-layer feed-forward neural networks as classifiers to distinguish between accreted and in situ stars. Neural networks are chosen over other methods, such as boosted decision trees, for two reasons. The first is the general increase in performance. The second is the ability to freeze part of the network in transfer learning, which will be discussed in great detail in Sect. 7. In the bulk of this paper, we will explore the power and limitations of this approach using simulated stellar catalogs, while in Sect. 8 we apply these lessons to data from *Gaia* DR2.

4.1. Neural network architecture

Our neural networks are implemented and trained with the Keras package (Chollet 2015) using the TensorFlow backend (Abadi et al. 2015). All networks in this study are constructed with five layers (the input, three hidden layers, and the output). The networks take between 4 and 9 measured quantities per star as inputs – these variations will be discussed in detail below. The hidden layers consist of 100 nodes each, using a ReLU activation function, that is, $h(x) = \max(0, x)$. The final output layer consists of a single node with sigmoid activation in order to scale the output to lie in the range from 0 to 1.

This network architecture was chosen as the result of a scan on a small set of the simulated data. The width of the layers, and the number of layers was allowed to change. The performance saturated when using three hidden layers of 100 nodes.

The neural network output S (star) denotes the value returned when applied to a given star, whose ideal behavior is

$$S(\text{star}) = \begin{cases} 1 & \text{accreted} \\ 0 & \text{in situ} \end{cases}. \quad (3)$$

In practice, a continuous range of outputs is produced; larger (smaller) values indicate greater network confidence that a star is accreted (in situ). The desired balance of background rejection versus signal efficiency translates into a choice of what value of $S(\text{star})$ to use as a selection cut.

Training is performed by adjusting the neural network weights to minimize a loss function, that is, a function chosen to smoothly vanish when truth labels and network output for a star agree, evaluated on a training set. Here, we use a weighted version of a standard choice for binary classification, the binary cross entropy:

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^N w_i (y_i \log f_i + (1 - y_i) \log (1 - f_i)), \quad (4)$$

where y_i is the truth label of the star (0 for in situ, 1 for accreted), f_i is the network prediction, N is the total number of stars in the sample, and w_i is the sample weight of the star (defined below).

4.2. Training and Error Sampling

To develop a robust classifier, we incorporate all uncertainties in the input variables into our training methodology. This avoids letting the classifier learn to heavily rely on kinematic or photometric properties of stars that are not actually well-measured. Additional justification for our error sampling approach is provided in Appendix B.

Because the in situ stars outnumber the accreted stars by a factor of ~ 100 for stars with $\delta\varpi/\varpi < 0.10$, we introduce compensating weights for the two populations⁵. A larger weight for

⁵ If the full catalog were used, the ratio would be smaller, especially at larger distances.

Table 2. Inputs used for star classifications.

Name	Explanation	Unit	Mean	σ
l	Galactic longitude	[deg]	2.72	78.2
b	Galactic latitude	[deg]	0.71	26.1
pmra	Proper motion in right ascension	[mas/year]	-1.66	10.6
pmdec	Proper motion in declination	[mas/year]	-2.83	11.0
parallax	Parallax	[mas]	0.48	0.86
phot_g_mean_mag	Extincted apparent G -band mean magnitude	[mag]	18.49	2.0
phot_bp_mean_mag	Extincted apparent G_{B_p} -band mean magnitude	[mag]	19.16	2.3
phot_rp_mean_mag	Extincted apparent G_{R_p} -band mean magnitude	[mag]	17.76	1.9
radial_velocity	Line-of-sight velocity	[km s ⁻¹]	-6.15	75.5
feh	[Fe/H]	[dex]	-0.20	0.42

Notes. The mean and standard deviation come from the m12i LSR0 mock catalog derived from the FIRE simulation and are used to rescale the measurements such that for m12i LSR0 all network inputs have a mean of zero and unit variance. These scale factors are applied to all the datasets used in this work: m12i, m12f, and *Gaia* DR2.

accreted stars ensures that the network is more sensitive to their properties during training. However, too large a weight would risk making the algorithms too sensitive to statistical fluctuations in the accreted subsample. Balancing between these concerns, we introduce a compensating factor of 5 in the weights relative to their proportion in the training set,

$$\begin{aligned} w_{\text{accreted}} &= \frac{1}{5} \frac{N_{\text{total}}}{N_{\text{accreted}}}, \\ w_{\text{insitu}} &= \frac{N_{\text{total}}}{N_{\text{insitu}}}. \end{aligned} \quad (5)$$

In addition, we set the size of the training batches so that there is an average of 5 accreted stars in each batch. This leads to a batch size of 915 for data sets using v_{los} measurements, and one of 565 when v_{los} is not used. We do not perform a comprehensive optimization over these choices, since we confirmed that the results are not very sensitive to the specific choice of relative weights or batch size. However, we note that the Keras default batch size of 32 yielded very inefficient training.

The training proceeds as follows. A given mock catalog is split into three subsets: one for training, one for validation, and one for testing. The validation and testing subsets each have a fixed size of 1 million unique stars. The training set is comprised of the remaining stars, in all cases consisting of at least 9.3 million stars.

For each star evaluated during training, we incorporate the impact of finite errors by generating 20 new instances of that star from a normal distribution whose center is the observed value and whose standard deviations are taken to be the observational uncertainties. The input variables for the stars (including all the “stars” generated by the random sampling procedure) are then rescaled so that the inputs are all roughly the same size. Specifically, the rescaling is chosen such that (over the stars of m12i LSR0 mock catalog with $\delta\varpi/\varpi < 0.10$) each input has a mean of 0 and a variance of 1. Table 2 provides the full list of inputs, as well as the mean and variance used for the rescaling. Note that this rescaling is performed with the values in Table 2 for all the datasets used in this work: m12i, m12f, and *Gaia* DR2. Before training, the network weights are initialized with the default glorot_uniform method (Glorot & Bengio 2010), which ensures that the initial variance of the inputs to each node is independent of its position in the network.

An epoch of training consists of two parts. First, we perform one iteration over all the batches in the training set, decreasing the loss function given in Eq. (4) for each batch in turn by updating the weights using the Adam optimizer (Kingma & Ba 2015). The initial learning rate is set to 10^{-3} . The learning rate controls how much the weights change. Then, after iterating over all of the batches in the training set, a similar procedure is performed using the validation set, including the random sampling over the observational errors. However, only the loss is calculated – the network weights are not updated during validation. This defines one epoch of training. The stars in the training set are then randomly shuffled and the next epoch begins.

If the loss calculated on the validation set does not improve for 5 epochs of training, the learning rate is reduced by a factor of 10, but is not allowed to decrease below 10^{-6} . Training ends when the validation loss has not improved for 10 epochs. This procedure typically takes ~ 50 epochs to complete. Specifying training completion in terms of validation set loss allows us to ensure that the network has the best potential to generalize to stars that it was not trained on. The value of the loss between the validation and training sets was comparable so we did not find it necessary to implement additional procedures that help avoid overfitting such as regularization or dropout. Once the network weights are fixed, all performance metrics are computed on the test set.

5. Tactics for hunting accreted stars

In this section, we explain our approach for optimizing the neural network configuration. Through extensive testing on the mock catalogs, we address the question of how many input measurements (and what error tolerances) are required to effectively yield a large-statistic, high-purity sample of accreted stars. First, we will explore the impact of eliminating the network’s access to one or more of the six phase-space degrees of freedom (d.o.f.). An optimal selection of inputs is critical because *Gaia* does not measure all six d.o.f. with the same level of accuracy. Then, we will investigate to what extent including photometry and/or metallicity into the analysis improves the network’s ability to discriminate. Finally, we will discuss how well the network generalizes from one viewpoint of a simulated galaxy to another. This informs how well we can expect a network trained on simulation to behave when applied to the actual *Gaia* dataset.

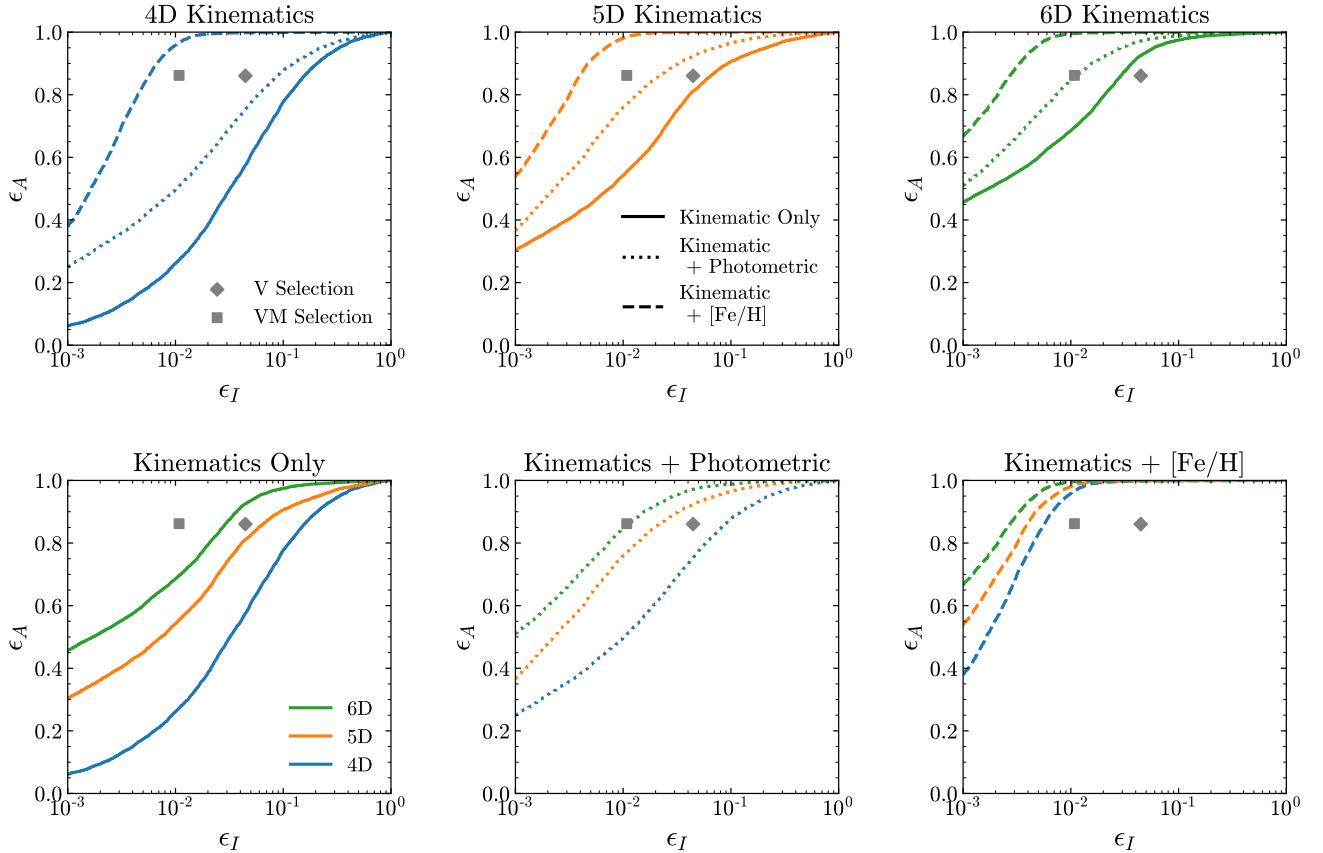


Fig. 3. Metrics on the test set of stars from the m12i LSR0 FIRE mock catalog with v_{los} measurements and $\delta\varpi/\varpi < 0.10$. Such ROC curves compare the accreted and in situ efficiencies, as defined in Eq. (2), for different networks. In general, a network has better performance if the accreted efficiency is high, and the in situ efficiency is low, that is, if the ROC curve tends to the top left of a given plot. In this figure, the label on each curve denotes what information the network had access to at training, and the plot title indicates the common network input for all curves in that plot. Note that the same nine curves are plotted in the *top row* and the *bottom row* – they are only organized differently to make the comparisons as transparent as possible. The gray symbols denote the performance for the V and VM selections described in Sect. 3. ZM is out of the range of the plots, as it has an in situ efficiency of 1.8×10^{-4} and an accreted efficiency of 3.3×10^{-2} . The network performance consistently beats the traditional methods in all cases where metallicity is included as an input, even if only 4D kinematics are available. The addition of photometric inputs provides a relative improvement compared to kinematics-only networks, although it is not as powerful as metallicity.

Further aspects of generalization to other simulations and *Gaia* data itself are discussed in Sect. 7.

5.1. Performance versus input dimensionality

We begin by considering how performance is affected when greater or fewer input phase-space d.o.f. are provided to the network. We use three choices of kinematic networks:

- 4D: 1, b, pmra, pmdec
 - 5D: 1, b, pmra, pmdec, parallax
 - 6D: 1, b, pmra, pmdec, parallax, radial_velocity
- In addition, we also test the impact of providing the networks with access to photometry or metallicity:
- Kinematic + Photometric: 1, b, pmra, pmdec, (parallax, [radial_velocity]), phot_g_mean_mag, phot_bp_mean_mag, phot_rp_mean_mag
 - Kinematic + [Fe/H]: 1, b, pmra, pmdec, (parallax, [radial_velocity]), feh
- where by (parallax, [radial_velocity]) we denote that the kinematics will be specified as 4D, 5D, or 6D (as defined above) in the results.

To compare networks, we use the receiver operating characteristic (ROC) curve. A ROC curve is a parametric curve generated by comparing the accreted (signal) versus in situ

(background) efficiency as the cut on the network output, which measures the network’s confidence in its classification, is scanned from 0 to 1. Note that for a given background rejection factor (ϵ_I^{-1}), a larger accreted efficiency ϵ_A represents a better classifier, so that an ideal classifier has a ROC curve that asymptotes to the top left corner of the plots, e.g., see Fig. 3. Multiple summary statistics for binary classifiers are available including area under the curve, F_1 score, Matthews correlation coefficient and others. Especially in imbalanced classification problems such as this, relying on a single metric combines differing sources of misclassification that can require markedly different treatments in subsequent physics analyses. We therefore present full ROC curves almost exclusively throughout this paper.

We compare the ROCs for various networks to the performance of the traditional selection methods introduced in Sect. 3. In this way, we can assess what input variables are needed to match or exceed their performance. As these traditional methods define a unique signal region, they do not yield a continuum of operating points. Therefore, each is represented as a single point on a ROC plot. Due to the requirements of most of these traditional methods, we are only able to apply them to a subset of the mock stars from the m12i LSR0 catalog: those with parallax measurement errors $\delta\varpi/\varpi < 0.10$ and radial velocities

v_{los} . Therefore, all of the results (both machine learning and cut-based) in this sub-section are performed using only this 6D subset.

Table 3 gives the division of the mock data into training, validation, and test sets. We train a series of neural networks following the procedure of Sect. 4.2, so that performance using the validation set are what determines when training is completed. Then ROC curves are computed using the test set. The resulting behavior is shown in Fig. 3 for the different networks. The upper panels show the results for networks that see 4D (on-the-sky position and velocity), 5D (adding parallax), or 6D (adding radial velocity) kinematics, from left to right. Within each panel, the solid, dotted, and dashed lines represent networks that use only kinematics, kinematics + photometry, or kinematics + metallicity, respectively. The grey markers display the working points for the cut-based methods. (The ZM selection lies outside the plot range, but is given in Table 1.)

Using 4D kinematics alone (Fig. 3, upper left panel), the neural network does not achieve the level of classification obtained by the traditional methods. However, by adding in the photometric inputs, the performance dramatically improves because the photometric magnitude is correlated with the distance (or parallax) and the metallicity. Furthermore, since most of the accreted stars in m12i have lower metallicity than the in situ stars, having direct access to [Fe/H] yields even better network performance.

Next, we train a network to use 5D kinematic inputs, where the parallax measurement is now made available. The bottom left panel of Fig. 3 shows that the 5D kinematics-only network using parallax has significantly better performance than the 4D kinematics-only case. As shown in the upper middle panel of Fig. 3, including parallax in the kinematics-only network allows for performance close to that of the V selection. When the additional photometric information is added, the network can nearly achieve the performance of the VM selection. As parallax and photometric information is available for a large fraction of the stars in *Gaia* DR2, this initially appears to be a very encouraging result. However, as we demonstrate below, networks that utilize photometric information do not generalize as well, and so the primary focus of this paper will be on the kinematic networks.

Finally, the right panel in the upper row of Fig. 3 includes the radial velocity information, giving the network access to the complete 6D phase space⁶. Unsurprisingly, having 6D information allows the networks to achieve excellent distinguishing power between accreted and in situ stars. When combined with photometric or metallicity information, the resultant networks outperform the standard methods. In all cases the networks provide superior performance when given access to equivalent information. Unfortunately, line-of-sight velocity measurements are currently only available for a small fraction of stars in *Gaia* DR2, see Table 8. However, these results certainly motivate updating our analysis for future *Gaia* data releases (as well as ground-based spectroscopic surveys such as LAMOST (Deng et al. 2012), DESI (DESI Collaboration et al. 2016), and SDSS-V (Kollmeier et al. 2017) that will include more stars with 6D measurements.

⁶ It would also be interesting to examine other combinations of input variables. For example, as parallax becomes harder to measure beyond ~ 4 kpc, a 5D combination of *l*, *b*, *pmra*, *pmdec*, and *radial_velocity* could allow for access to stars that are farther away. Going to a 3D model with *l*, *b*, and *radial_velocity* could extend the reach in distance even further. However, the method chosen here is applicable to a larger number of stars, so we leave an exploration of these other choices for future studies aimed at subsequent *Gaia* releases and spectral surveys.

Table 3. Detailed composition of stars within the m12i LSR0 dataset with a measurement of radial velocity v_{los} and parallax error $\delta\varpi/\varpi < 0.10$.

$\delta\varpi/\varpi < 0.10; v_{\text{los}}$	N_{accreted}	N_{insitu}	Purity
Training set	51 209	9 349 637	0.55%
Validation set	5364	994 636	0.54%
Test set	5462	994 538	0.55%

Notes. The stars are divided into training, validation, and testing sets when utilized for training the neural network and testing its subsequent performance. For each set, we provide the number of true accreted and in situ stars as well as the purity, which is defined as the number of truly accreted stars selected divided by the total number of stars in the particular dataset.

The bottom row of panels in Fig. 3 contain the same nine curves, but organized by the extra information provided to the network: kinematics only, kinematics + photometric, and kinematics + metallicity, respectively. The different curves now indicate the cases of 4D, 5D or 6D kinematics. While adding more information is clearly better, this makes it clear that systematically there is a larger gain in going from 4D to 5D than from 5D to 6D, especially when only the kinematic information is provided (left bottom panel).

These results demonstrate that deep learning can exploit hidden correlations in the data to identify accreted stars. That being said, it is not obvious how this approach can generalize if the test/validation sets include incomplete information about the stars, or when applying the networks trained on one location to another viewpoint within a simulated galaxy. We turn to addressing these questions next.

5.2. Generalizing from brighter to dimmer stars

The previous subsection focused on the subset of stars with small parallax errors and radial velocities. We now consider how the network performance is affected if we train a network on the subset of stars that are closest to the viewpoint (and hence satisfy the requirement of having $\delta\varpi/\varpi < 0.10$ and a measurement of v_{los}) using only 5D kinematic information, and then apply the network to a larger data set consisting of stars that are farther away and therefore are not measured as well. Comparing the star counts in Table 3 and 4 reveals that the size of the mock dataset increases by about a factor of 5 by relaxing these quality cuts. However, this increase is not evenly distributed between accreted and in situ stars; the fraction of truth-level accreted stars within the training set increases from 0.55% to 0.88%. We will see that indeed the network is learning general enough features for this extrapolation to yield useful results, which will ultimately bolster our ability to produce a high statistics sample of accreted stars.

In Sect. 5.1, we demonstrated the extent to which using 5D rather than 4D kinematics improves network classification performance. Repeating this test on the mock data with no v_{los} requirement, the conclusions are unchanged. Therefore, we only show the results for 5D kinematics going forward.

The ROC curves for the 5D networks are presented in Fig. 4. The blue lines correspond to the networks that were trained on stars with $\delta\varpi/\varpi < 0.10$. In contrast, the networks depicted by the orange lines were trained on stars that were additionally required to have a line-of-sight velocity measurement, that is, training only involves the more nearby stars. Note that these

Table 4. Same as Table 3, but with no requirement of v_{los} , that is, requiring parallax errors $\delta\varpi/\varpi < 0.10$ only.

$\delta\varpi/\varpi < 0.10$	Accreted	in situ	Purity
Training set	430 376	48 266 382	0.88%
Validation set	8938	991 062	0.89%
Test set	8828	991 172	0.88%

Notes. Removing the requirement on v_{los} increases the sample size by a factor of ~ 5 and the number of accreted stars by a factor of ~ 9 .

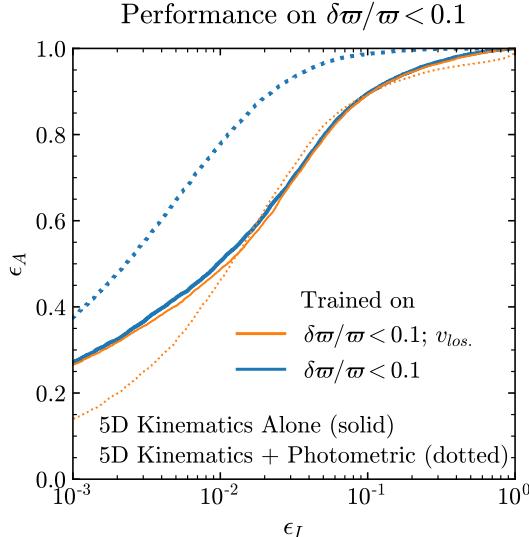


Fig. 4. Testing how networks trained on nearby or bright stars generalized to farther or dimmer stars, using test data of the m12i LSR0 FIRE mock catalog. We compare results where the network is trained on stars with $\delta\varpi/\varpi < 0.1$ with either v_{los} measurements required (orange) or not (blue). Solid or dotted lines indicate if the network uses 5D kinematics or also includes photometry as inputs. All networks are tested on the 5D dataset, such that stars have a small parallax error, but they may or may not have a radial velocity measurement. The network that only uses 5D kinematics gives equivalent results regardless of whether it is trained on data with radial velocities or not, that is, the solid blue and orange lines are comparable. However, we find that when 5D kinematics + photometry are used as inputs, the network performance is significantly hampered when training on the data set with radial velocities and testing on the broader data set with no v_{los} requirement, that is, the dotted orange line is suppressed relative to the dotted blue line.

results are different than those presented in Fig. 3, because the network is now applied to test stars which may or may not have line-of-sight velocities.

First, we highlight the two solid lines in Fig. 4, which only use 5D kinematic information as network inputs. The orange and blue solid lines are essentially indistinguishable. This indicates that there is little drop in performance when the requirement of a radial velocity measurement is imposed, which makes sense when the network is only trained using 5D kinematics. Looking forward, this makes it plausible that we can use the small subset of stars within *Gaia* DR2 that have radial velocity measurements to train a network on actual stars. If the network only uses the 5D kinematic information for the stars, the network will safely generalize to stars with no radial velocity measurements.

Next, we consider what happens when photometry is included in the training (dotted lines in Fig. 4). When training and testing on stars that may not have radial velocities (blue), the inclusion of photometric data improves the performance of

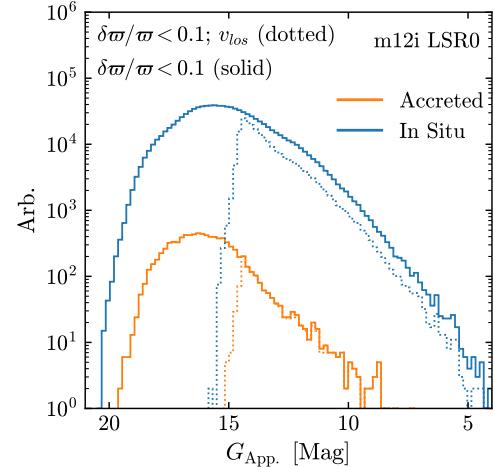


Fig. 5. Distributions of the mean apparent brightnesses in the G band for Tables 3 and 4. The normalization of the distribution is in arbitrary units. In situ stars are indicated by the blue lines, and accreted stars are shown as orange lines. Networks that are trained on stars with $\delta\varpi/\varpi < 0.10$ and v_{los} measurements (dotted lines), using both 5D kinematics and photometry as inputs, do not perform well when tested on all stars with $\delta\varpi/\varpi < 0.10$ (solid lines). This result was demonstrated in Fig. 4, and here we expose the root of the problem. A network that is trained on stars where v_{los} is required only sees relatively bright stars and does not generalize well when applied to a data set where dimmer stars are present.

the network, as anticipated from the previous section. However, something surprising happens when the network is trained on stars with radial velocity measurements and is then applied to the less restrictive data set (orange line). Contrary to expectations, this network performs the worst when high purity is demanded. The explanation for this can be inferred from the apparent brightness plots for the two mock data sets in Fig. 5. This figure shows that only bright stars have a line-of-sight velocity measurement. Since the network that is trained on a more restrictive dataset and does not see stars with an apparent magnitude above 15, it does not properly identify the dimmer stars in the expanded data set.

Additionally, we also computed the ROC curves for stars in the test set with $\delta\varpi/\varpi < 0.1$ and $G_{\text{App.}} < 15$ but no restriction on v_{los} . There are roughly twice as many stars in this set than when requiring v_{los} , and the results generalize better between training on stars with or without v_{los} requirements. However, this only applies to less than half of the stars with $\delta\varpi/\varpi < 0.1$.

We will eventually use transfer learning to train the network on stars in the Milky Way using empirical labels which require spectroscopy. This motivates us to focus our attention on the 5D kinematic network. Going forward, the networks will not be given access to photometric information. Appendix C shows the resulting analysis when using the network with photometry.

The results of this section have demonstrated that it is possible to use the measurements available for a large portion of *Gaia* DR2 to separate in situ and accreted stars. Next, we will explore the extent to which a network trained on one location within a galaxy can be applied to another.

5.3. Exploring multiple LSRs

Three synthetic *Gaia* surveys are provided for m12i (Sanderson et al. 2020), each separated in the plane of the disk by angular intervals of $2\pi/3$. These surveys allow one to study the effects of features like the bar and spiral arms, which break rotational symmetry. Moving between these viewpoints

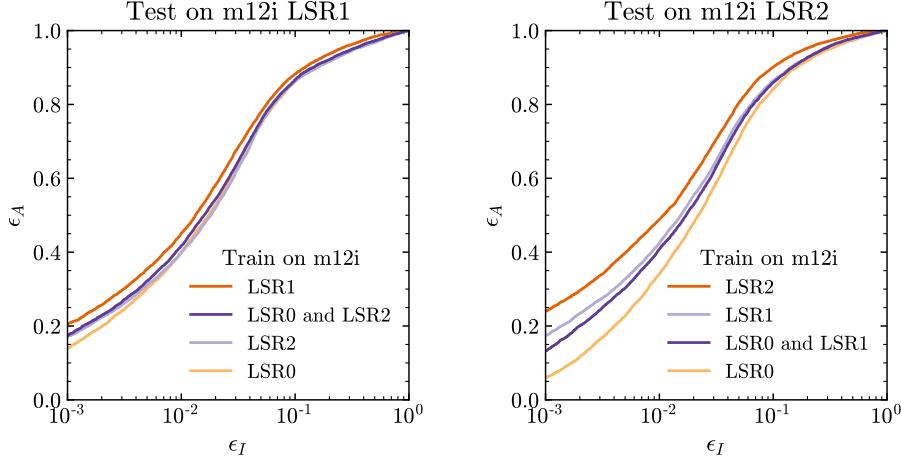


Fig. 6. ROC curves showing the performance of networks that are trained on one local standard of rest (LSR) of `m12i` and applied to a different LSR of `m12i`. The networks are trained and tested on the subset of stars that have small parallax errors, and only use 5D kinematics as inputs. The orange line shows the result of training and testing on the same LSR; this provides a benchmark for the best that the network can do. In cases where the network is trained on a different LSR than which it is tested, the performance is not dramatically worsened. This suggests that the network is learning gross features that distinguish accreted from in situ stars, rather than detailed features of a specific galaxy.

provides a way to validate the level to which the networks rely on features that are specific to a particular LSR. One might be concerned that the neural network is performing so well by memorizing localized substructures. Since moving between LSRs causes these substructures to occupy very different regions of the 5D phase space (and possibly removes them), passing this test tells us that our algorithms are learning gross features of the galaxy, as opposed to merely remembering the particulars of a specific viewpoint.

Thus far, all results have utilized mock data taken from `m12i` LSR0. Now we will assess the impact of moving to a different position using the surveys of `m12i` denoted by LSR1 and LSR2. We again select stars with 5D kinematics (small parallax errors and no requirement on line-of-sight velocity measurements). From this subset, we then assign 1 million stars for validation and 1 million for testing. For reference, we note that the LSR1 and LSR2 datasets are slightly larger than that of LSR0; the stars in each survey are unique.

We train new neural networks on each of the LSRs. The plots in Fig. 6 show the resulting ROC curves derived from a number of different networks that are only trained on the 5D kinematic information for each of the three LSRs and a combination of the two not being used for testing. Each network is then tested on the LSR1 and LSR2 datasets (left and right panels, respectively). In both panels, the orange line (top in the legend) shows the performance that can be achieved using each local training set and applying it to the same LSR. This serves to benchmark the performance of the various networks as it allows our networks to see local kinematic information during training. The yellow lines (bottom in the legend) depict the networks that are trained on LSR0 and applied to either LSR1 (left panel) or LSR2 (right panel); the lavender lines in the left (right) panel depict the networks trained on LSR2 and applied to LSR1 and vice versa, respectively. In both cases, the lavender ROC curves out-perform the yellow, showing that LSR1 and LSR2 are more similar to each other than to LSR0.

The first lesson we learn from these figures is that the networks are relying on distinguishing features that are universal to all three LSRs, which is why even the worst curve (in yellow) still does a good job rejecting background. For the LSR0 curve, at a fixed false positive rate ϵ_I of 0.01, the accreted efficiency ϵ_A

is 0.40 (0.34) for the left (right) plot. As a comparison, the VM selection had a false positive rate of 0.017 and selected accreted stars at a rate of 0.89 – but it requires both 3D velocities and metallicity, which are not necessarily present in this data set. On the other hand, the loss of performance experienced by the networks trained on stars from the same catalog (the orange lines) show that the networks are taking advantage of some additional detailed structures that does not appear in all frames of reference.

Since we want to keep our neural network from learning too many specific details of any particular simulated reference frame, we choose to train it on multiple mock surveys simultaneously. The purple curves in Fig. 6 display the result of this test. In the left panel, we see that the network trained on both LSR0 and LSR2 behaves similarly to that trained only on LSR2, which is better than the one only using LSR0. In the right panel, we see that the combined network does not perform as well as that using LSR1, but does do better than that using LSR0. However, in both cases the performance of the combined network is closer to the higher-performing viewpoint extrapolation. This is encouraging, since *a priori* we have no way of knowing which choice of training sample will give better performance. In comparison to training on the worse-performing LSR0 alone, at a false positive rate of 0.01 the combined network gives a gain in the accreted efficiency of 9% (19%) for the left (right) plot.

We conclude that training on multiple mock catalogs both raises the baseline performance and helps the networks to focus on general features to distinguish accreted from in situ stars without overemphasizing particular local substructure. Two comments are in order. First, we remark that all of the examples explored here were taken from the same simulated galaxy. One might be concerned that the network is still learning specifics of `m12i`, which are common among the different LSR catalogs (even though the stars themselves only fall into one catalog, since we have restricted the distance from the mock satellite to $\lesssim 4$ kpc). As part of Sect. 7, we address this issue by utilizing another simulated galaxy (`m12f`) and by applying a method known as transfer learning. The second comment is that there is a downside to having an overly general network, since the ideal machine would know about detailed properties of the Milky Way. We at least partially accommodate this desire by bootstrapping our training with actual Milky Way data in order to let the

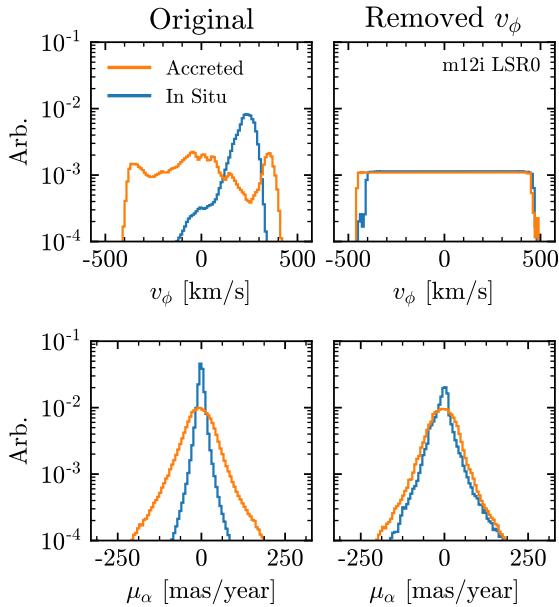


Fig. 7. Data planing allows one to determine which high-level variables play an important role in the network classification. In this procedure, one starts with the initial distribution (in arbitrary units) of a particular variable, in this example, the rotational velocity v_ϕ (*top left*), and then suppresses the information by weighting each star inversely by the probability that it falls within a particular histogram bin (*top right*). The planing procedure effectively removes the v_ϕ information from the dataset, but also affects the distribution of other variables. For example, the *bottom row* shows the proper motion in the right ascension, μ_α , distributions before and after planing in v_ϕ . By running the networks on the planed datasets, one can estimate the importance of the planed variable in driving the classification. The distributions in this figure pertain to the subset of m12i LSR0 with small parallax errors and a measurement of v_{los} .

machine learn and subsequently expose the local substructure present in the *Gaia* data.

6. What the machine is using to classify

The results presented thus far make it clear that the neural network can distinguish between accreted and in situ stars. This section presents an analysis that utilizes the “data planing” method (Chang et al. 2018) to suggest what information is important for classification.

The power of neural networks is that they can learn complex representations of the data from low-level inputs. The goal of this technique is to expose the high-level variables being utilized by the machine. To this end, we construct a “planed” dataset by removing high-level information from the original data. One can then train a new network on the planed data, with the drop in resulting network performance as a quantitative measure of the importance of that variable for the original classifier. An example of a high-level variable that will prove to carry a lot of discriminating power is the rotational velocity of the stars in the galactic frame. We note that the network must learn something analogous to both variable estimation and coordinate transformation on the 5D kinematic inputs to take advantage of this information, since the velocity measurements provided are neither complete nor in cylindrical coordinates.

When planing the data, we suppress information by first generating a histogram in the variable of interest for the accreted and in situ stars separately. Each star is then given a weight inversely proportional to the probability that it falls in a particular bin.

Table 5. Area under the ROC curve (AUC) for networks trained on different planed datasets of the m12i LSR0 catalog, and using only 5D kinematics.

Data	AUC
Using full set	0.96
Removed ϕ	0.97
Removed R	0.94
Removed z	0.90
Removed v_z	0.84
Removed v_R	0.83
Removed v_ϕ	0.75

Notes. The extent to which the AUC is lowered when removing information represents how important that variable is in distinguishing accreted stars. For comparison, we also provide the AUC for the original (un-planed) data. Of the simple kinematic variables considered, the disk rotation is the most crucial, lowering the AUC from 0.96 to 0.75. This suggests that, even though the machine is not given enough information to fully reconstruct v_ϕ , it is still inferring a set of correlations of the input variables that correspond to this quantity.

These weights are applied to each sample within the loss function during training, which is analogous to uniformly sampling over the given variable. Therefore, when we plane in, e.g., the rotational velocity, the network should not see the peak associated with the stars rotating with the disk.

Any observable that is correlated with the planed variable will have an altered distribution with respect to the original data set. Figure 7 shows an example of the planing process on m12i LSR0 for the stars with $\delta\varpi/\varpi < 0.10$ and a v_{los} measurement. The left column displays the initial distributions of v_ϕ and the proper motion in the right ascension direction (μ_α) on the top and bottom panel, respectively. The distributions after weighting the samples in inverse proportion to the v_ϕ probabilities are shown in the right column. By design, the v_ϕ distributions in the planed data set are uniform, up to statistical noise. The lower right plot for μ_α shows that uniformly sampling the stars over v_ϕ impacts the distributions of the input variables.

Planing the data results in poorer network performance. To quantify this, we use the area under the ROC curve (AUC) as a simple indicator of a network’s ability to separate signal from background. An AUC of 1 implies that every star is perfectly identified, while an AUC of 0.5 is equivalent to a random guess. By comparing the results of networks trained on different planed datasets, we can assess how important different observables are for the classification. The more important the planed variable, the larger the reduction of the AUC.

The AUC of networks trained on different planed datasets is shown in Table 5. We choose to use Galactocentric cylindrical coordinates (v_R, v_ϕ, v_z) as the high-level variables, as this reflects the symmetry of the disk. Using the full dataset (without planing), the base AUC score is 0.96. Removing the ϕ positional information of the stars actually slightly improves the score, presumably because this transformation of the data simply enforces the underlying cylindrical symmetry of the disk. Within the solar neighborhood, we also do not expect the galactic distance to play a large role, and we see that removing this only slightly changes the performance. For contrast, since the in situ stars are clustered around $z = 0$ kpc (the disc), planing in z reveals that this variable is somewhat important to the network. Interestingly, even though the network has access to the 3D position and only 2D velocity, evenly sampling over any of the velocity components reduces the performance more than removing position. As anticipated,

this shows quantitatively that the rotation of the disk is the most significant variable for classifying the accreted stars.

Even though the networks are not using the information shown in Table 5 directly, these results hint at what the machine is learning. Specifically, there is more discriminating power in the kinematic distributions than in the position information. This additionally explains how the networks are able to generalize so well when testing on different LSRs in a specific galaxy: v_ϕ does not vary dramatically among them. The dominant importance of information about velocities within the plane of the galactic disk is precisely what one would expect from our current understanding of galactic dynamics.

We note that while v_ϕ is the most important variable considered here, it does not reduce the AUC to 0.5, reflecting that the machine is relying on additional information. Correlations among the variables considered here are likely important. However, planing is difficult in multiple dimensions, especially when the underlying distributions are not smooth, and is beyond the scope of this work.

7. Transfer learning on FIRE

In the last section, we showed that neural networks are able to take advantage of non-trivial kinematic correlations among the 5D inputs to distinguish stars accreted onto a galaxy versus those born within the galaxy. For instance, even though the algorithm is not given v_ϕ directly – there is in fact not even enough information available to fully reconstruct v_ϕ – it still infers that the combination of observables roughly corresponding to v_ϕ is an important quantity. However, all of these results are based on a study of a single simulated galaxy, and so one may be concerned that the kinematic features being used to discriminate will not be relevant for other galaxies. In this section, we will demonstrate that this worry is unfounded, in that the network’s performance is largely maintained when applying it to a different galaxy with a distinct merger history. To account for the differences between galaxies, we will use of a scheme known as “transfer learning” to refine a network’s performance on the particular galaxy of interest. Our implementation of this approach requires initial training on a simulated galaxy, followed by additional restricted training on (simulated or real) data from another galaxy.

7.1. Transfer learning methodology

Transfer learning refers to a training strategy whose goal is to utilize a model trained for one task (perhaps where there is a large high quality data set) to perform a different, albeit related, task (Caruana 1994; Bengio 2011; Bengio et al. 2011; Donahue et al. 2013; Yosinski et al. 2014). The utility of this approach has been demonstrated for image recognition tasks, where the idea is that certain layers of a deep network have learned to recognize specific features in an image, for example, eyes or legs. However, a network that was trained on many images of high-quality stock photos may not work well for blurry, poorly-lit user images, even though the objects to be identified have many features in common with the stock photos. Thus, instead of initializing the “blurry” network with random weights, the paradigm of transfer learning starts with the weights determined by training on the stock photos. This procedure is often referred to in the literature as “pre-training.” Ideally, this implies that the network starts its second round of training near the minimum of the loss function, such that only minor tweaks are needed to optimize for the new data.

Our network will be pre-trained on simulated galaxy catalogs (in analogy with the stock photos), which are similar to

the Milky Way. By training on multiple catalogs, we will obtain a network that is only sensitive to general dynamical features. Then, transfer learning will be performed on a separate data set from another simulated galaxy. This will result in a network that can accurately classify accreted and *in situ* stars between two galaxies with different merger histories.

7.2. Transfer learning experiments

To choose the best method of transfer learning for our particular application, we perform the following set of numerical experiments using **m12i** as the basis for the pre-training, and **m12f** as the mock “real” galaxy. As we have chosen to use 5-layer networks, there are multiple ways in which the weights of the different layers can be unfrozen to allow the network to specialize. We consider the following possibilities:

1. Update only the first layer. This assumes that the generic network has already learned most of the relevant features, but that the input normalizations vary from catalog to catalog. Note that before feeding the data to the networks, the inputs have been normalized by the mean and variance of the **m12i** LSR0 data. By re-training the weights of the first hidden layer, the network can learn the normalizations specific to another galaxy.
2. Update only the last layer. This option still assumes that all of the learned features are essentially the same, but that the final combination of these features that ultimately yields the optimal classification is different for the new galaxy. Re-training the network fixes the combination of high-level features learned in the pre-trained network.
3. Update the first and last layers. We also consider a combination of options two and three.
4. No transfer learning. As a means of comparison, we provide results where the pre-trained network is used directly without updating any of the weights.

There is an additional aspect of the method that requires testing. The stars in the new galaxy that are used for re-training must be given a label. While we have access to truth-level information in the simulations, we obviously do not for the Milky Way. Therefore, we also test which of the traditional cut-based methods works best as a way to derive such labels. This allows us to evaluate whether it is better for the network to see more examples of truly accreted stars (while inherently mislabelling more *in situ* stars as accreted), or to use stars that have a much higher probability of being accreted (while incorrectly labeling many accreted stars as *in situ*).

We take the **m12f** LSR0 catalog to be our “real” data set. The initial pre-training uses a combined dataset of the LSR0, LSR1, and LSR2 catalogs of **m12i** for stars with $\delta\varpi/\varpi < 0.10$, providing us with an initial training set containing 176 842 422 *in situ* stars and 1 427 323 accreted stars. In the transfer learning step, we use 200 000 randomly drawn stars from the **m12f** LSR0 catalog with $\delta\varpi/\varpi < 0.10$ and measured radial velocities⁷. The initial learning rate for the transfer learning training step is set to 10^{-4} . Although the results are similar when using a larger initial learning rate, the overall performance suffers. The resulting networks derived for all the different transfer learning options are compared against a test set of 10 000 000 stars pulled from the larger **m12f** LSR0 dataset, that is, these stars are not required to have a line-of-sight velocity measurement.

⁷ This sample size was chosen to be comparable to the RAVE DR5-*Gaia* DR2 catalog (see Table 8), which is the data set that will be used when ultimately performing transfer learning for the Milky Way.

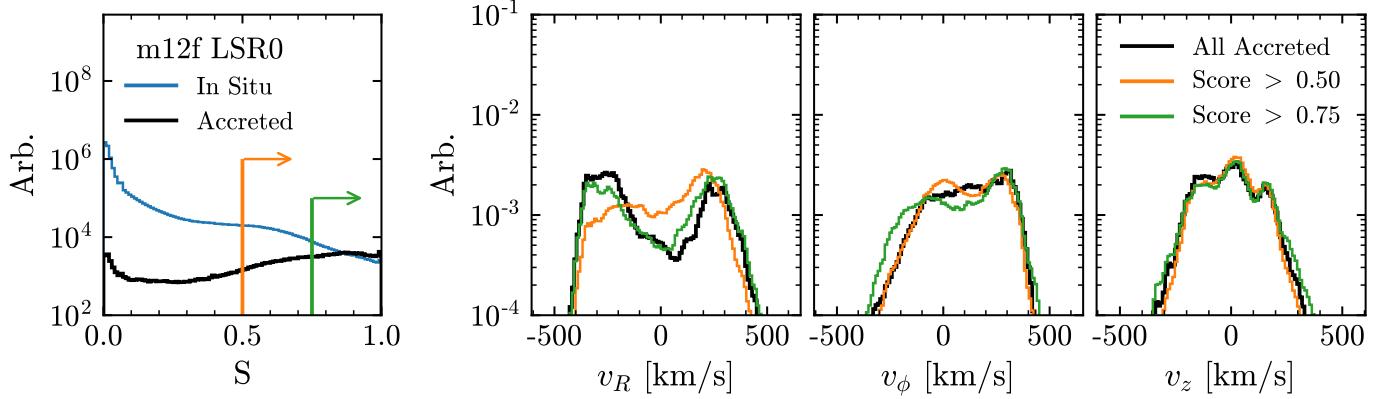


Fig. 8. Example metrics for transfer learning. *Left panel:* distribution in arbitrary units of the network scores for the true accreted (black) and in situ (blue) test stars of m12f LSR0. We note that the stars that are truly in situ have network scores peaked towards 0, while those that are truly accreted have scores peaked towards 1. The orange and green arrows indicate two cuts on the network scores: $S > 0.50$ (orange) and $S > 0.75$ (green). For the example illustrated here, the scores are specific to a network where transfer learning was performed on the last layer using the ZM selection to derive labels. In the remaining panels, the orange and green lines show the normalized v_R , v_ϕ , and v_z distributions of the stars with scores larger than the indicated cut. The thick black lines correspond to the distribution of the truth-level accreted stars, not just the ones passing the cuts. We see that cutting on a network score of 0.75 (green lines) better reproduces the truth distributions. To quantify the goodness-of-fit, we calculate the χ^2 for the v_R , v_ϕ , and v_z distributions separately, and sum them together to get a total $\sum \chi^2$. The lower the value of $\sum \chi^2$, the better the network reproduces the truth distributions. For the case illustrated here, the $\sum \chi^2$ is a factor of 3 smaller for the green distributions, compared to the orange ones.

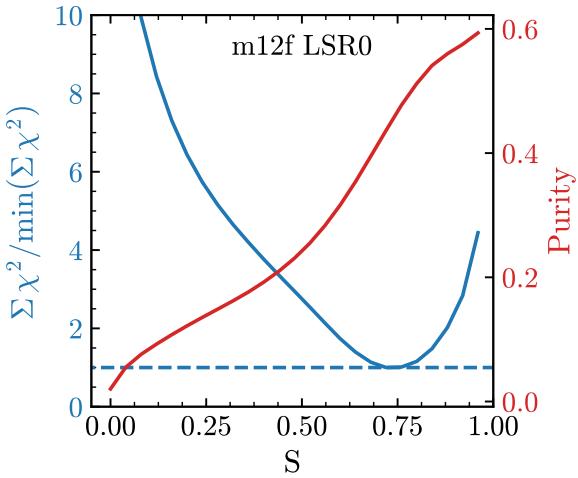


Fig. 9. Determination of the transfer learning method. The blue curve shows the sum over the χ^2 of the normalized v_R , v_ϕ , and v_z distributions compared to all of the accreted stars in the test sample, normalized to the minimum value obtained. The x axis shows the cut value, where all stars with a network output larger than this are selected. The network was trained with transfer learning updating only the last layer and used the ZM labels. A cut of 0.75 gives the best overall fit. The red curve shows the purity of the resulting sample (and corresponds to the labels on the right axis). Tighter cuts on S result in a more pure sample, but may bias the resulting kinematic distributions when stronger than $S \sim 0.75$.

We need a metric to evaluate which transfer learning approach is optimal. To this end, we compare the recovered distributions for the 3D velocities (v_R , v_ϕ , v_z) to the truth-level distributions. The left-most panel of Fig. 8 shows the network score for the truth-level accreted (black) and in situ stars (blue) for one example. In this case, the training labels for the m12f LSR0 set used for transfer learning were determined using the ZM selection and transfer learning was only performed on the last layer of the network. As desired, the network scores are peaked towards 1 for the truly accreted stars, whereas the in situ stars have scores peaked towards 0. The remaining panels in Fig. 8 show

the normalized velocity distributions, where the thick-black line is the truth-level distribution for the accreted stars. The orange and green lines show all of the stars from the test set (not just the accreted stars) that have network scores greater than 0.5 and 0.75, respectively. The orange distributions do not match those of the accreted stars, yielding χ^2 values of 2.46×10^{-2} , 2.12×10^{-3} and 1.54×10^{-3} for v_R , v_ϕ , and v_z , respectively, where χ^2 is defined as

$$\chi^2 = \sum_{i \in \text{bins}} (f(v)_i - h(v)_i)^2, \quad (6)$$

and $f(v)_i$ is the value of the distribution of the truly accreted stars in velocity v bin i and $h(v)_i$ is the same except for the stars passing the neural network cut. In comparison, the χ^2 values for the green distributions are 4.38×10^{-3} , 4.77×10^{-3} , and 1.07×10^{-3} . A cut on the network score of 0.75 yields a lower sum of these χ^2 values, $\sum \chi^2$, by about a factor of 3. This suggests that cutting on a network score of 0.75 yields distributions that are closer to truth; this is clear from Fig. 8 where the green lines trace the black better than the orange lines do.

This procedure suggests an operative definition of an optimal value for the cut on S , as shown in Fig. 9. The blue line corresponds to the left axis, and shows $\sum \chi^2$ relative to the minimum value. The red curve (correspond to the right axis) shows the purity of the sample after making the cut. A cut of 0.75 minimizes $\sum \chi^2$, while a tighter cut results in a more pure sample.

We repeat this procedure for each transfer learning and empirical labeling method, finding the optimal cut value on the network score using $\sum \chi^2$ in the same way for each case. The optimal cut value typically falls between 0.5 and 0.9 for each method. This optimal cut is not always close to unity, since although increasing the value of the cut increases the fraction of stars that are accreted in the final selection, doing so also biases the distributions. In other words, the selection of stars that are most easily identified as accreted do not reproduce the underlying full distribution of accreted stars.

Table 6 shows the benchmark performance for the different networks tested here through their minimum resulting $\sum \chi^2$

Table 6. Results for a neural network trained on LSR0, LSR1, and LSR2 of m12i using only 5D kinematics as inputs.

Trainable layers		$\sum \chi^2 [\times 10^{-2}]$
No transfer learning		1.07
VM	1st	2.13
	Last	1.20
	1st and last	2.28
ZM	1st	1.04
	Last	1.02
	1st and last	1.03

Notes. We compare the results with and without transfer learning on m12f LSR0. When the transfer learning is performed, either the first layer is updated, or the last layer is updated, or both. In the transfer learning step, all truth labels are derived using either the VM or ZM selections, as defined in Sect. 3. $\sum \chi^2$ quantifies how similar the v_R , v_ϕ , and v_z are to the truth-level distributions, as illustrated in Fig. 8. For each transfer learning method, we determine the optimal cut on the network score, as described in the text. In general, the ZM selection with transfer learning performed by only varying the last layer does the best job at reproducing the kinematic distributions of the true accreted stars, that is, it achieves the lowest $\sum \chi^2$.

values. Interestingly, using the VM selection as the transfer learning label always results in worse fits. This can be traced to the fact that the VM selection has large contamination from in situ stars, thereby negatively affecting the results. In addition, as this selection uses velocities (which as shown above are more important than spatial information), it has more potential to create a biased set. In contrast, the ZM selection label yields lower $\sum \chi^2$ values, likely because of the low false-positive rate of this method. As a result, we choose to use the ZM selection and only update the last layer the network⁸. Furthermore, we take the optimal cut on the neural network output to be 0.75 so as not to bias the velocity distributions – any star with a network output greater than this will be labeled as accreted.

7.3. Transfer learning validation

The methodological decisions in our transfer learning regimen and the determination of the optimal cut value on the neural network output were justified by pre-training networks on m12i and then transferring and testing on m12f LSR0. The goal of this section is to validate these choices. In particular, we will apply the machinery of the previous section to a different viewpoint within m12f, specifically LSR1. Truth-level information will only be utilized for computing performance metrics.

As before, we begin with the networks pre-trained on m12i. Next, we mock the dataset used for the transfer learning training step by randomly selecting 200 000 stars from the m12f LSR1 catalog with the requirement that they have $\delta\varpi/\varpi < 0.10$ and a radial velocity measurement. Stars within this subset that have $[\text{Fe}/\text{H}] < -1.5$ & $|z| > 1.5$ kpc are assigned a label of 1, and the stars that do not pass this cut are labeled with 0. The inputs are rescaled according to Table 2. The last network layer is then unfrozen and transfer training is performed.

⁸ It could be argued that performing transfer learning to get a 5% better fit is unnecessary, especially when it relies on using empirical labels to define “truth”. However, we will find that the use of transfer learning is a critical step towards obtaining a reliable catalog of real Milky Way stars. This will be discussed more in Sect. 8 where the actual catalog is presented.

The results are presented in Fig. 10 for all test stars in m12f LSR1 with $\delta\varpi/\varpi < 0.10$. Note that we are able to plot the 6D information since we have access to the truth-level phase space. The blue lines are the distributions of the accreted stars in the m12i LSR0 catalog, depicting the shapes on which the network was pre-trained. The orange lines show the truth-level distributions for the accreted stars in m12f LSR1; these clearly differ in important ways from the m12i LSR0 distributions. The black, dotted lines display the stars from the m12f LSR1 catalog selected after the transfer learning. Importantly, the distributions of the stars that are identified as accreted by the network closely trace the truth-level distributions of m12f LSR1. This demonstrates that the network after transfer learning can successfully select accreted stars from a galaxy that is not the same as the one used in pre-training, even though the distributions of the two galaxies have distinct features. It is a little troubling that the network does classify some younger stars as accreted; however, this is a small fraction of the selected stars. As the network is only using kinematic information, it is hard to distinguish what is causing this selection.

Of the 10 million stars in the test set, only 163 727 are accreted. When using the cut of 0.75, there are 187 721 stars selected, of which 77 614 are accreted. The purity of the selection is 41%, comparable to the standard methods in Sect. 3, while the accreted selection efficiency is 47%. The cut value can be increased to get a more pure sample at the risk of biasing the distributions. If the cut is placed at 0.95 then there are only 35 495 stars selected, with 20 897 being accreted. This brings the purity to 59%. While this results in a selection efficiency of only 13%, this is better than any of the methods in Table 1. These results are summarized in Table 7.

This analysis has demonstrated that it is possible to classify accreted versus in situ stars using simulated data that has been processed to very closely resemble *Gaia* DR2. The optimized strategy relies on beginning with a network originally trained on simulated galaxies followed by the use of transfer learning to adjust the network by augmented training on real data. We conclude that it is reasonable to apply this method to the *Gaia* data, where the transfer training data set is taken from stars in *Gaia* DR2 that have been cross matched with spectroscopic surveys. Cutting on the neural network output (at the optimized value) yields a *Gaia* DR2 catalog of accreted stars with reasonable purity and unbiased kinematic distributions. Table 7 acts as an estimate of the performance we expect in the *Gaia* catalog.

Finally, we note that with the large sample sizes considered here, there are examples of accreted stars that the network scores very low and also in situ stars that the network gives high scores. Unsurprisingly, we find that the accreted stars that the network thinks are in situ tend to have kinematics very similar to that of the disk. The in situ stars that get high scores tend to be either at larger distances from the galactic plane, or have large retrograde velocities.

8. Deriving the catalog

Having optimized our approach using mock catalogs, we turn to constructing neural networks that will be used to distinguish accreted from in situ stars within the real *Gaia* DR2 data. The number of stars in the *Gaia* DR2 release passing the various cuts are shown in Table 8. The first stage of training utilizes the combined datasets for the three LSR catalogs of m12i, and takes around 26 h on dual E5-2690v4 (28 cores) nodes. Then transfer learning is implemented by updating the weights of the last layer of the network by training on the cross-matched

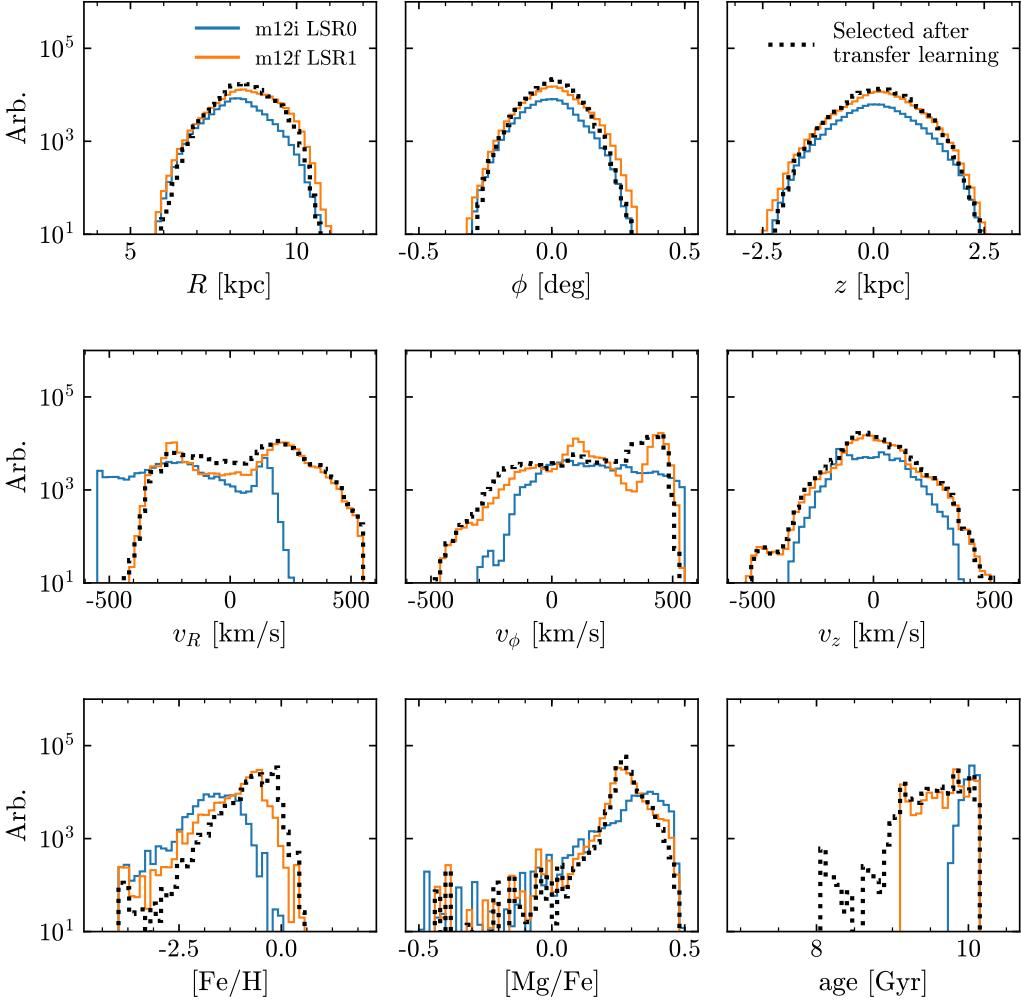


Fig. 10. Validation of the transfer learning procedure. All distributions are shown in arbitrary units. The network is pre-trained on m12i LSR0 using 5D kinematics as inputs. The true 6D distributions of accreted stars that the network sees are depicted by the blue lines in each panel. Then, transfer learning is performed by updating the last layer of the network on m12f LSR1 using 200 000 stars that are labeled using the ZM selection; that is, they are labeled as accreted if $|z| > 1.5$ kpc and $[\text{Fe}/\text{H}] < -1.5$. Stars from m12f LSR1 that have a network score greater than 0.75 are then marked as accreted. Their distributions are indicated by the black dotted lines in the panels. For comparison, we also show the truth-level distributions for accreted stars in m12f LSR1 as the orange lines. The distributions of the stars selected by the network do an excellent job at reproducing the truth-level distributions for the accreted stars in m12f, even though the network was pre-trained on an entirely different galaxy with a different merger history. This result justifies our confidence that a network trained on simulation can be successfully applied to the Milky Way.

Table 7. Test set of the m12f LSR1 catalog using stars with $\delta\varpi/\varpi < 0.10$.

Data set	Purity	ϵ_A
$S(\text{star}) > 0.75$	41%	47%
$S(\text{star}) > 0.95$	59%	13%

Notes. The neural networks are trained on separate m12i galaxy simulations with truth-level information and then re-trained using only observable information from m12f LSR1. The optimal cut of 0.75 yields velocity distributions most similar to the truth-level accreted distributions, but only has a purity of 41%. This is comparable to the traditional methods listed in Table 1. A stronger cut of 0.95 increases the purity, but decreases the fraction of accreted stars that are selected, and can bias the velocity distributions. This provides an estimate for the purity and efficiency expected when applied to *Gaia*.

RAVE DR5-*Gaia* DR2 catalog (Kunder et al. 2017) and takes around 1 h. There are 347 979 stars in the cross-matched catalog

that have $\delta\varpi/\varpi < 0.1$, a measurement of the line-of-sight velocity, and have iron abundance. We label any star that further passes a $|z| > 1.5$ kpc and $[\text{Fe}/\text{H}] < -1.5$ cut as a 1 for training; this cut yields a total of 849 stars. Stars that do not pass this cut are labeled as 0. The distribution in the $b-l$ plane for both sets of stars are shown in the top panel of Fig. 11⁹.

Some resulting distributions for this training data are shown in the bottom panel of Fig. 11. Note that while the stars given a training label of 1 are relatively evenly distributed in velocity space (the left and middle panels), they are tightly clustered in the HR diagram (the lower right panel). We do not take extinction or reddening into account in the HR diagram (Gaia Collaboration et al. 2018b). Additionally, the v_R-v_ϕ plot makes it clear that some of the training stars that are labeled

⁹ We verified that restricting the training stars to follow a similar distribution on the sky does not yield any important effects on the transfer learning. We repeated the analysis of Sect. 7.3 using a selection function similar to RAVE, while still only selecting 200 000 stars. The final results only changed minimally compared to those of Fig. 10.

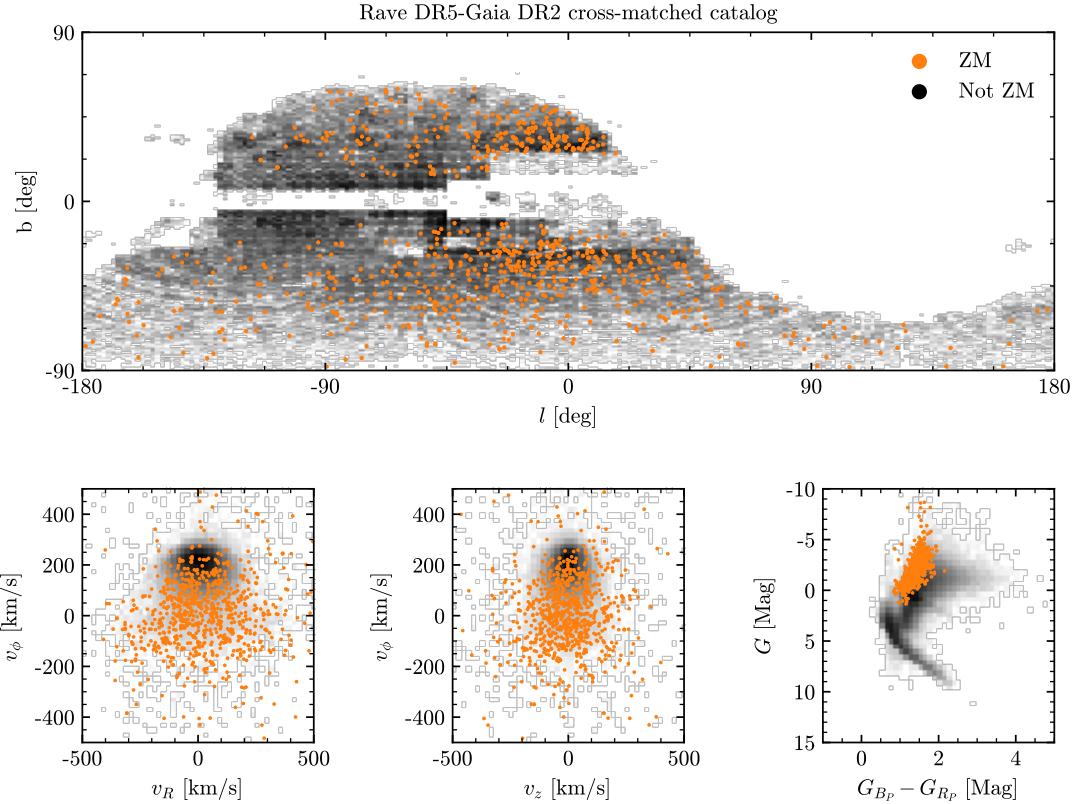


Fig. 11. Stars used for transfer learning to re-train the networks that were pre-trained on the `m12i` catalogs from FIRE. The points marked in orange have $|z| > 1.5$ kpc and $[\text{Fe}/\text{H}] < -1.5$ and are labeled as accreted for the training. The remaining stars are marked as in situ during the training step. Due to the large sample size, the stars labeled as in situ are displayed as a density plot, where the darker bins have a larger number of stars. The HR diagram in the lower right panel does not take into account the extinction and reddening (Gaia Collaboration et al. 2018b).

Table 8. Number of stars in the *Gaia* DR2 dataset.

<i>Gaia</i> DR2	Number of stars
Full set	1 692 919 135
$\delta\varpi/\varpi < 0.10$	72 488 013
$\delta\varpi/\varpi < 0.10; v_{\text{los}}$	5 393 495
$\delta\varpi/\varpi < 0.10; \text{RAVE DR5}$	347,979

Notes. Requiring a line-of-sight velocity measurement or metallicity (for example as a cross match with RAVE DR5) greatly reduces the number of available stars.

as accreted have rotations that are consistent with the disk, even though they lie outside of the thin disk (by construction). This is important because it shows the network that accreted stars in the Milky Way can have disk-like rotations.

It is worth noting that the testing on mock data presented in Sect. 7 did not make an overwhelmingly compelling case that transfer learning is necessary. However, transfer learning has a much bigger impact when applied to *Gaia* DR2. As a cross-check, we performed the same analysis only using the pre-trained network, without transfer learning on the *Gaia* DR2 stars. Only 0.1% of the stars were selected as accreted, while the simulations suggest between 0.89–1.6% of the stars with $\delta\varpi/\varpi < 0.10$ should be accreted. With the transfer learning, we select 1.06% of the stars. We interpret this result as implying that transfer learning allows the neural network to “unlearn” some hidden correlations among the star kinematics within the simulated datasets. One such potential effect is that each star particle

in the simulation is used as a seed for many stars in the catalog, which could yield some unphysical correlations.

8.1. Validating on milky way data

We perform one final validation of our approach, the results of which are presented in Fig. 12. First, we split the cross-matched RAVE-*Gaia* dataset into thirds. Then we treat one of these thirds of the data as a test set, and we perform the transfer learning by training on the other two thirds. The resulting network is applied to the test set. We repeat this procedure twice more by taking the pre-trained network and treating each of the other thirds as the test set. This allows us to maximize the statistics for validation while ensuring that the network has not seen any of the test stars when training.

As justified above, we classify any star with a neural network output score above 0.75 as accreted, while any star with a lower score is in situ. Some properties of this final validation are provided in Fig. 12, which shows the resulting 1D distributions along the diagonal, with the stars labeled as in situ or accreted in blue or orange, respectively. Additionally, we show the distributions of the stars labeled by applying the ZM selection cuts. The in situ stars exhibit a distinct peak in the v_ϕ distribution corresponding to the rotation of the disk. The stars labeled as accreted are roughly symmetric about 0 km s^{-1} in all three velocity components, as expected. Finally, the $[\text{Fe}/\text{H}]$ distribution shows that the in situ stars tend to have larger metallicities than the accreted stars. Recall that the network is not provided with any metallicity information as an input. The metallicity distribution of the accreted stars identified by the network extends across the full observed range.

Rave DR5-Gaia DR2 5D Kinematic Network Selection

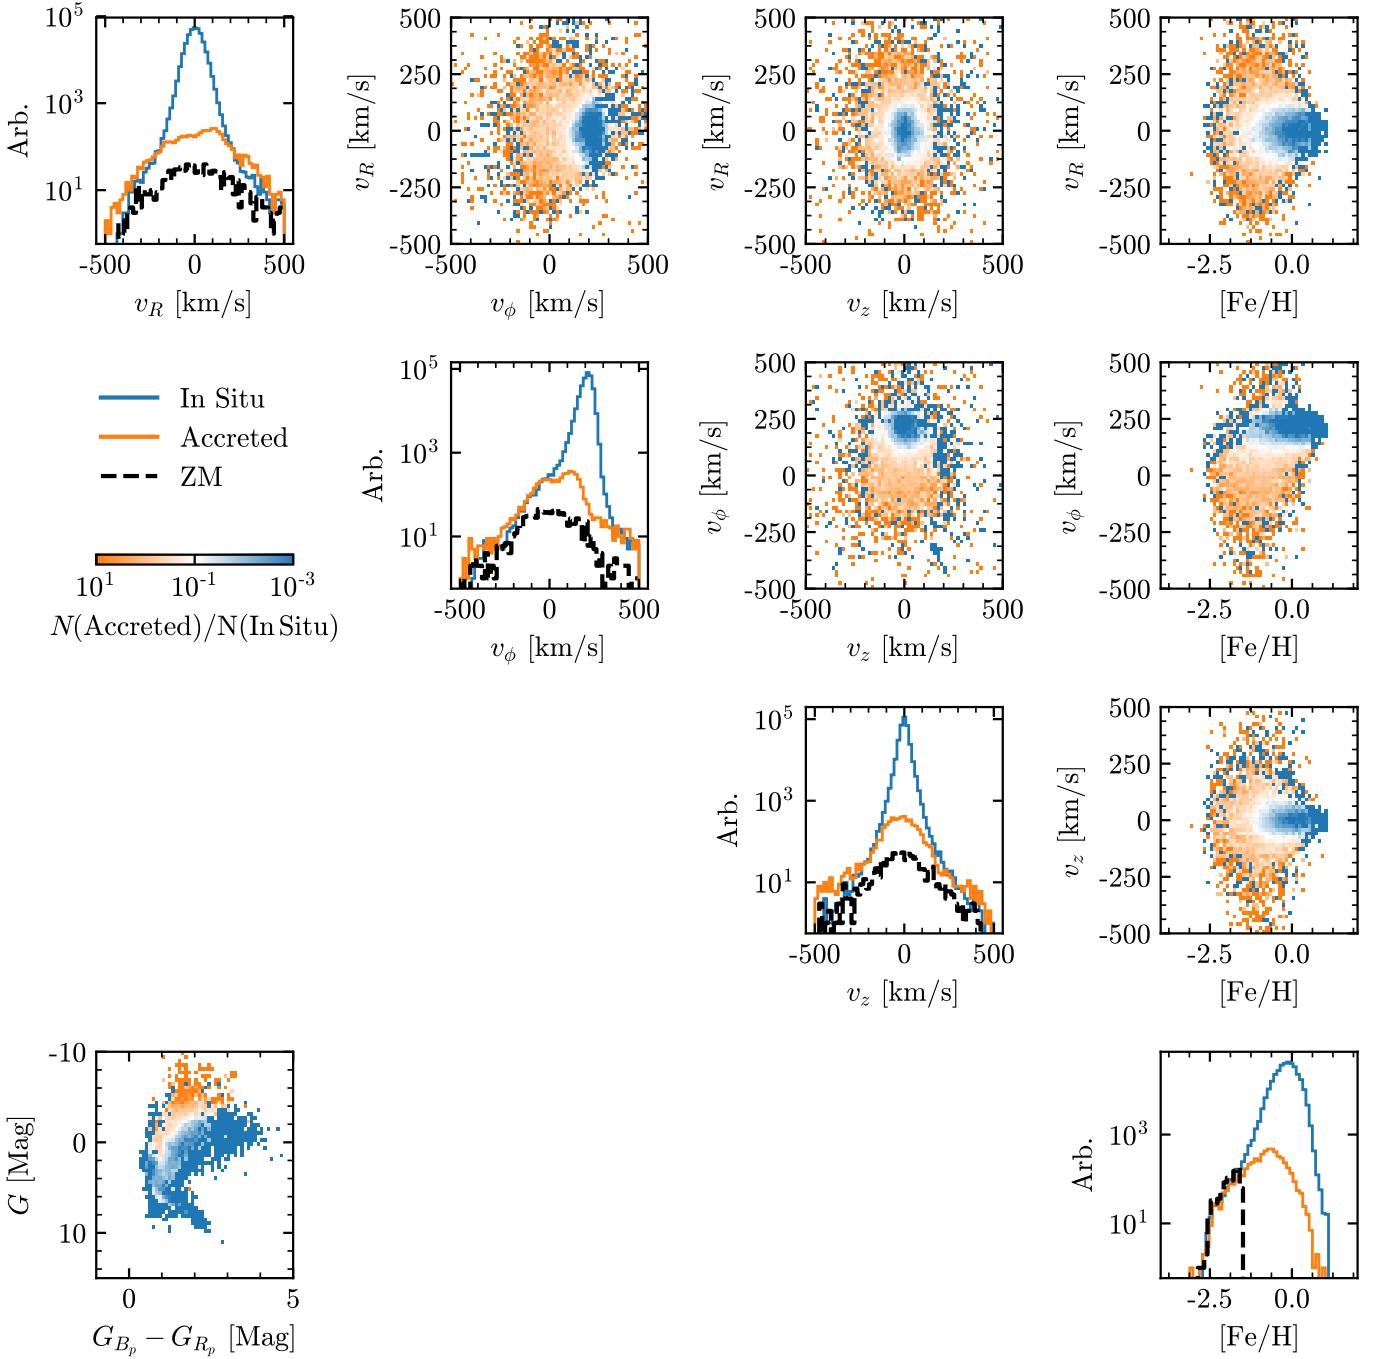


Fig. 12. Distributions of the stars in the RAVE DR5–Gaia DR2 catalog. The blue and orange lines denote stars labeled as in situ or accreted by the network using only 5D kinematics. The black-dashed lines show the distributions of the stars selected with the ZM selection method. The stars selected by the network have similar kinematic profiles, but extend to larger metallicities. The fact that these distributions agree with the expected halo distributions gives us confidence in our ability to derive a catalog of accreted stars for which we do not have access to a metallicity measurement.

The remaining panels in Fig. 12 show 2D histograms of the ratio of the number of accreted stars to the number of in situ stars per bin. The bluest bins have a factor of 10^3 more in situ stars than accreted (or they contain no accreted stars). The orange bins have more accreted stars than in situ stars, where the largest ratio observed in any given bin is 10. These distributions reveal that

the in situ stars dominate specific regions of phase space, while the accreted stars span a much larger range.

We interpret Fig. 12 as providing further validation that the kinematic network is working as expected. This provides confidence that our *Gaia* DR2 catalog of accreted stars will be physically meaningful.

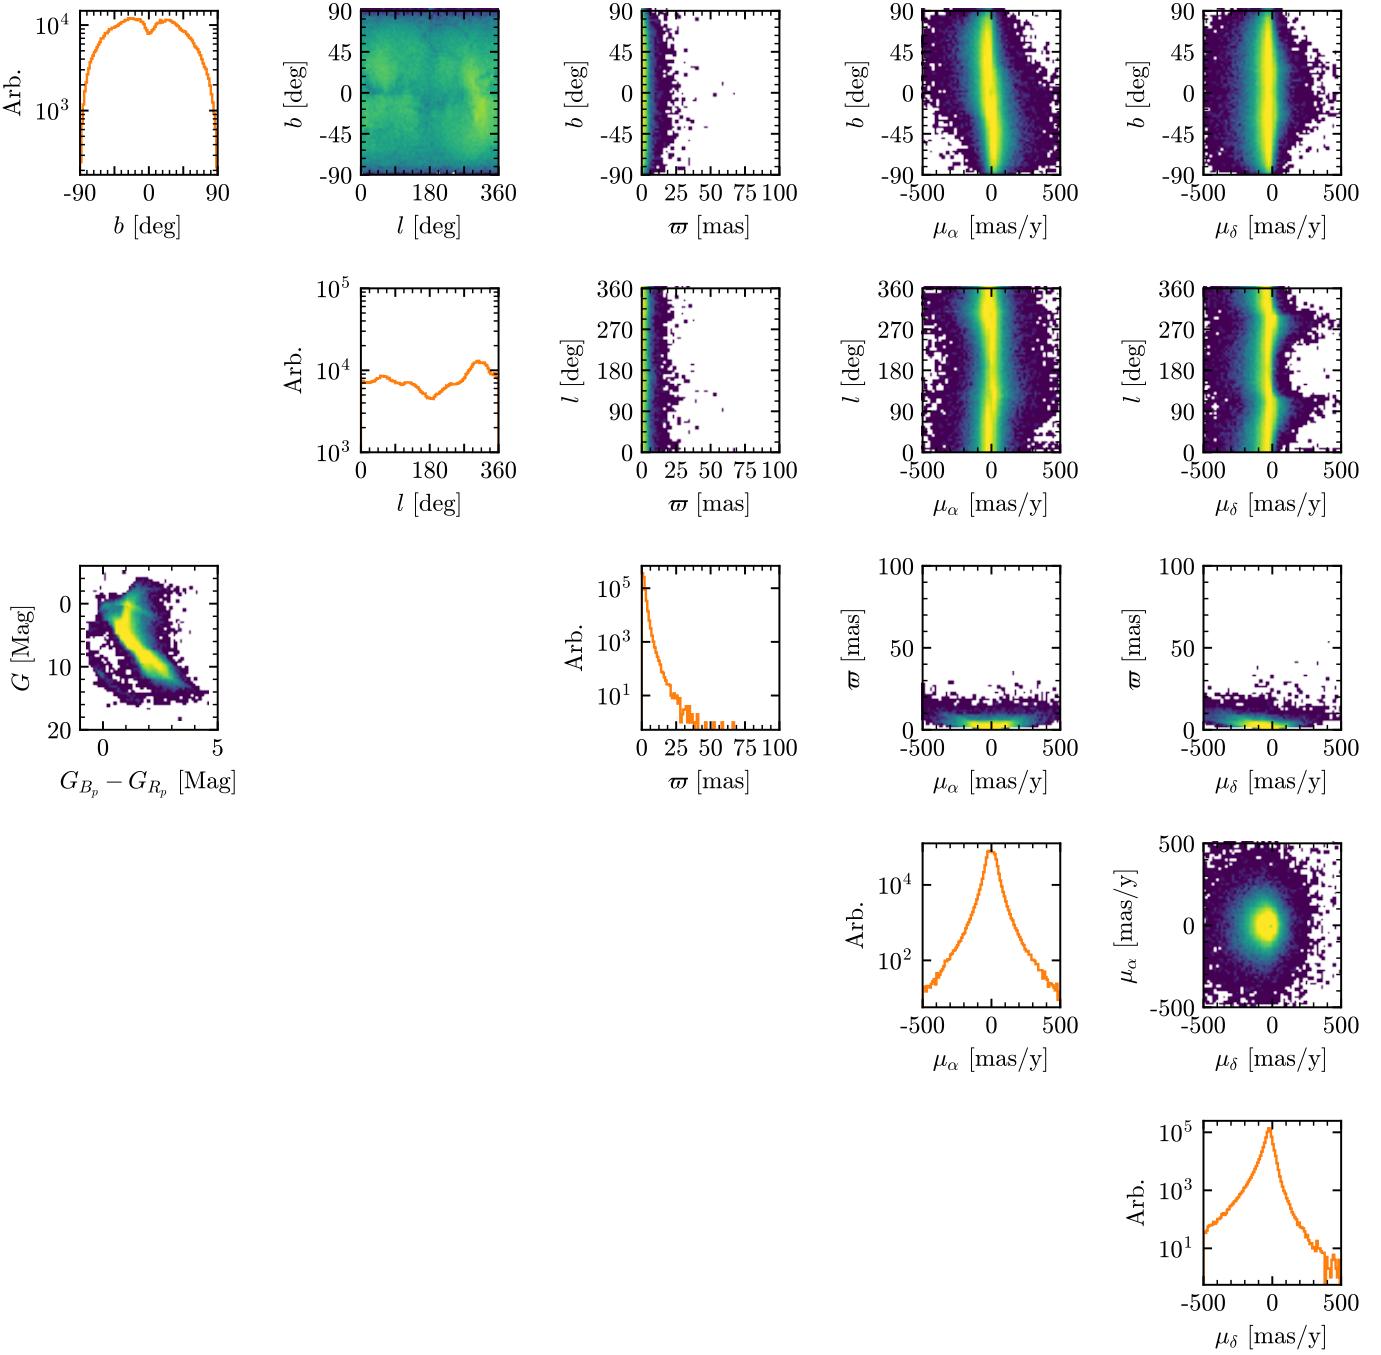


Fig. 13. Stars within *Gaia* DR2 with $\delta\varpi/\varpi < 0.10$. These stars have been identified as accreted by our neural network using 5D kinematics alone. While the HR diagram is shown, photometric information is not used in these selections. As expected for an accreted population, these stars are nearly uniformly distributed in the l, b plane. See Fig. A.1 for the analogous figure for the stars classified as in situ.

8.2. A Catalog of accreted stars

Now we are finally ready to obtain the neural network used to generate the catalog. We again start by pre-training on simulation (the mix of all three LSRs within m12i with $\delta\varpi/\varpi < 0.10$ and no v_{los} requirements), and then perform transfer learning using the full RAVE DR5–*Gaia* DR2 cross-matched catalog as the training data, with labels determined by the ZM selection. Since we do not separate off a test set, we are not able to perform any validation tests of the final network. However, this choice has the benefit that it allows us to maximize our statistics during transfer

learning to obtain the best possible performance for our ultimate *Gaia* DR2 catalog.

After this last stage of training, the network is applied to every star within the *Gaia* DR2 catalog with $\delta\varpi/\varpi < 0.10$, which yields ~ 72 million total stars, as shown in Table 8. Stars with a network score greater than 0.75 are called accreted and added to our catalog. This yields a total of 766 993 stars that are identified by the neural network as being accreted. The distributions are shown in Fig. 13, and the resulting catalog is available on Zenodo (Ostdiek et al. 2019). Figure A.1 provides the distributions of the stars marked as in situ for comparison.

9. Conclusions

In this work, we generate a catalog of stars identified by a neural network as having been accreted onto the Milky Way, as opposed to forming *in situ*. Extensive testing was provided to justify a methodology based on transfer learning methods. One possible limitation of this study is the necessity of training on computationally expensive simulations. The pre-training was performed on a mixture of simulated mock *Gaia* data derived for three different viewpoints within the m12i galaxy from the FIRE simulations. The last layer of this network was then unfrozen and allowed to re-train on RAVE–*Gaia* cross-matched data, where the training labels were derived with a traditional cut-based selection (specifically the ZM selection). The resulting network was applied to the subset of *Gaia* DR2 stars with small parallax errors, $\delta\varpi/\varpi < 0.10$. Stars with a neural network score greater than 0.75 are tagged as accreted. The resultant catalog is available online¹⁰.

The testing performed above used mock *Gaia* catalogues generated from cosmological zoom-in hydrodynamic simulations. Within the simulations, we are able to derive “truth-level” labels for the stars using the merger tree history. We first explored how much information as measured by *Gaia* was necessary to accurately classify stars. This was a critical test since not every star in *Gaia* DR2 has been observed accurately enough to provide a complete measurement of its 6D phase space with small errors. From these explorations, it was found that using 5D inputs – l , b , parallax, $pmra$, and $pmdec$ – both provides good classification accuracy and the ability to generalize from nearby stars with full 3D velocity measurements to those that are further away such that v_{los} is unknown. In addition, we showed how transfer learning can be used to allow a network trained on one simulation to generalize to an independent simulation, using a small amount of local data for re-training. This local data does not need to have truth-level information, which would be unavailable for a real galaxy, but could be labeled using traditional cut-based techniques to identify accreted stars. We note that these traditional techniques can only be used for a small subset of the stars, since they often rely on either knowledge of the full 6D phase space, or spectroscopic measurements that require cross matching between *Gaia* DR2 and another dataset. We showed that using a more restrictive definition of accreted stars for labeling the stars during the transfer training worked better than a more inclusive definition, which justified our choice of the ZM selection (the requirement that $|z| > 1.5 \text{ kpc}$ and $[\text{Fe}/\text{H}] < -1.5$) for this purpose.

After identifying the specifics of our transfer learning methodology, we took the network pre-trained on the simulated galaxy viewpoints, and re-trained the network on the RAVE DR5–*Gaia* DR2 cross-matched catalog. The resulting network was then applied to the 72 million stars in the *Gaia* DR2 catalog that have parallax measurement errors of less than 10 percent. Of these, our network identifies 766 993 stars as having accreted onto the Milky Way. Because the network does not require full 3D velocity information, the number of identified stars is a factor of 21 larger than what would be possible using traditional (3D velocity-dependent) methods. This large increase in the number of stars allows for a much more detailed study of the halo structure and the history of the Milky Way. In a companion paper, (Necib et al. 2019a), the catalog is used to reproduce known structures, such as *Gaia* Enceladus and the Helmi stream, and identify a new structure. The new structure is studied in more detail in another companion paper, (Necib et al. 2019b),

where we argue the new structure is the remnant of a merged dwarf galaxy. Our catalog is a new tool for performing Galactic archaeology that can be leveraged to unlock novel discovery potential thus far hidden within *Gaia* DR2.

Acknowledgements. We are grateful to Ben Farr and Graham Kribs for useful discussions. This work utilized the University of Oregon Talapas high performance computing cluster. BO and TC are supported by US Department of Energy (DOE), under grant number DE-SC0011640. LN is supported by the DOE under Award Number DE-SC0011632, and the Sherman Fairchild fellowship. MF is supported by the Zuckerman STEM Leadership Program and in part by the DOE under grant number DE-SC0011640. ML is supported by the DOE under contract DE-SC0007968 and the Cottrell Scholar Program through the Research Corporation for Science Advancement. AW is supported by NASA, through ATP grant 80NSSC18K1097 and HST grants GO-14734 and AR-15057 from STScI, and a Hellman Fellowship from UC Davis. SGK and PFH are supported by an Alfred P. Sloan Research Fellowship, NSF Collaborative Research Grant #1715847 and CAREER grant #1455342, and NASA grants NNX15AT06G, JPL 1589742, 17-ATP17-0214. Numerical simulations were run on the Caltech compute cluster “Wheeler”, allocations from XSEDE TG-AST130039 and PRAC NSF.1713353 supported by the NSF, and NASA HEC SMD-16-7592. This work was performed in part at Aspen Center for Physics, which is supported by National Science Foundation grant PHY-1607611. We also are grateful for the support from the 2018 CERN-Korea TH Institute. This research was supported by the Munich Institute for Astro- and Particle Physics (MIAPP) of the DFG cluster of excellence “Origin and Structure of the Universe”. This research was supported in part by the National Science Foundation under Grant No. NSF PHY-1748958. RES thanks Nick Carriero, Ian Fisk, and Dylan Simon of the Scientific Computing Core at the Flatiron Institute for their support of the infrastructure housing the synthetic surveys and simulations used for this work. This work has made use of data from the European Space Agency (ESA) mission *Gaia* (<http://www.cosmos.esa.int/gaia>), processed by the *Gaia* Data Processing and Analysis Consortium (DPAC, <http://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the *Gaia* Multilateral Agreement. Funding for RAVE has been provided by: the Australian Astronomical Observatory; the Leibniz-Institut fuer Astrophysik Potsdam (AIP); the Australian National University; the Australian Research Council; the French National Research Agency; the German Research Foundation (SPP 1177 and SFB 881); the European Research Council (ERC-StG 240271 Galactica); the Istituto Nazionale di Astrofisica at Padova; The Johns Hopkins University; the National Science Foundation of the USA (AST-0908326); the W. M. Keck foundation; the Macquarie University; the Netherlands Research School for Astronomy; the Natural Sciences and Engineering Research Council of Canada; the Slovenian Research Agency; the Swiss National Science Foundation; the Science & Technology Facilities Council of the UK; Opticon; Strasbourg Observatory; and the Universities of Groningen, Heidelberg and Sydney. The RAVE web site is at <https://www.rave-survey.org>.

References

- Abadi, M., Agarwal, A., Barham, P., et al. 2015, *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*, Software available from tensorflow.org, <https://www.tensorflow.org>
- Astropy Collaboration (Price-Whelan, A. M., et al.) 2018, *AJ*, **156**, 123
- Bahcall, J. N., & Soneira, R. M. 1980, *ApJS*, **44**, 73
- Behroozi, P. S., Wechsler, R. H., & Wu, H.-Y. 2013, *ApJ*, **762**, 109
- Belokurov, V., Evans, N. W., Irwin, M. J., et al. 2007, *ApJ*, **658**, 337
- Belokurov, V., Erkal, D., Evans, N. W., Koposov, S. E., & Deason, A. J. 2018, *MNRAS*, **478**, 611
- Belokurov, V., Zucker, D., Evans, N., et al. 2006, *ApJ*, **642**, L137
- Bengio, Y. 2011, *Proceedings of the 2011 International Conference on Unsupervised and Transfer Learning Workshop in - Volume 27, UTLW'11 (JMLR.org)*, 17
- Bengio, Y., Bastien, F., Bergeron, A., et al. 2011, in Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics, eds. G. Gordon, D. Dunson, & M. Dudík (Fort Lauderdale, FL, USA: PMLR), *Proc. Mach. Learn. Res.*, **15**, 164
- Bienayme, O., Robin, A. C., & Creze, M. 1987, *A&A*, **180**, A94
- Bland-Hawthorn, J., & Gerhard, O. 2016, *ARA&A*, **54**, 529
- Bonaca, A., Conroy, C., Wetzel, A., Hopkins, P. F., & Kereš, D. 2017, *ApJ*, **845**, 101
- Borsato, N. W., Martell, S. L., & Simpson, J. D. 2019, *MNRAS*, **492**, 1370

¹⁰ <https://doi.org/10.5281/zenodo.3354470>

- Bullock, J. S., & Johnston, K. V. 2005, [ApJ](#), **635**, 931
- Bullock, J. S., Kravtsov, A. V., & Weinberg, D. H. 2001, [ApJ](#), **548**, 33
- Carleo, G., Cirac, I., Cranmer, K., et al. 2019, [Rev. Mod. Phys.](#), **91**, 045002
- Caruana, R. 1994, [Proceedings of the 7th International Conference on Neural Information Processing Systems, NIPS'94](#) (Cambridge, MA, USA: MIT Press), 657
- Chang, S., Cohen, T., & Ostdiek, B. 2018, [Phys. Rev. D](#), **97**, 056009
- Chollet, F., et al. 2015, Keras. <https://keras.io>
- Chomiuk, L., & Povich, M. S. 2011, [AJ](#), **142**, 197
- Deng, L.-C., Newberg, H. J., Liu, C., et al. 2012, [Res. Astron. Astrophys.](#), **12**, 735
- DESI Collaboration (Aghamousa, A., et al.) 2016, ArXiv e-prints [arXiv:1611.00036]
- Domínguez Sánchez, H., Huertas-Company, M., Bernardi, M., et al. 2019, [MNRAS](#), **484**, 93
- Donahue, J., Jia, Y., Vinyals, O., et al. 2013, ArXiv e-prints [arXiv:1310.1531]
- Eggen, O. J., Lynden-Bell, D., & Sandage, A. R. 1962, [ApJ](#), **136**, 748
- Fattah, A., Navarro, J. F., Sawala, T., et al. 2016, [MNRAS](#), **457**, 844
- Faucher-Giguère, C.-A., Lidz, A., Zaldarriaga, M., & Hernquist, L. 2009, [ApJ](#), **703**, 1416
- Fernández-Trincado, J. G., Beers, T. C., Tang, B., et al. 2019a, [MNRAS](#), **488**, 2864
- Fernández-Trincado, J. G., Moreno, E., & Pérez-Villegas, A. 2019b, [MNRAS](#), submitted, [arXiv:1904.05370]
- Gaia Collaboration (Brown, A. G. A., et al.) 2018a, [A&A](#), **616**, A1
- Gaia Collaboration (Babusiaux, C., et al.) 2018b, [A&A](#), **616**, A10
- Garrison-Kimmel, S., Wetzel, A., Bullock, J. S., et al. 2017, [MNRAS](#), **471**, 1709
- Garrison-Kimmel, S., Hopkins, P. F., Wetzel, A., et al. 2018, [MNRAS](#), **481**, 4133
- Glorot, X., & Bengio, Y. 2010, [Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics](#), 249
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, [ApJ](#), **622**, 759
- Grand, R. J. J., Gómez, F. A., Marinacci, F., et al. 2017, [MNRAS](#), **467**, 179
- Grand, R. J. J., Helly, J., Fattah, A., et al. 2018, [MNRAS](#), **481**, 1726
- Helmi, A., & White, S. D. M. 1999, [MNRAS](#), **307**, 495
- Helmi, A., Veljanoski, J., Breddels, M. A., Tian, H., & Sales, L. V. 2017, [A&A](#), **598**, A58
- Helmi, A., Babusiaux, C., Koppelman, H. H., et al. 2018, [Nature](#), **563**, 85
- Herzog-Arbeitman, J., Lisanti, M., Madau, P., & Necib, L. 2018, [Phys. Rev. Lett.](#), **120**, 041102
- Hezaveh, Y. D., Perreault Levasseur, L., & Marshall, P. J. 2017, [Nature](#), **548**, 555
- Hopkins, P. F. 2015, [MNRAS](#), **450**, 53
- Hopkins, P. F., Narayanan, D., & Murray, N. 2013, [MNRAS](#), **4**, 432
- Hopkins, P. F., Wetzel, A., Kereš, D., et al. 2018a, [MNRAS](#), **480**, 800
- Hopkins, P. F., Wetzel, A., Kereš, D., et al. 2018b, [MNRAS](#), **477**, 1578
- Hopkins, P. F., Grudić, M. Y., Wetzel, A., et al. 2020, [MNRAS](#), **491**, 3702
- Huertas-Company, M., Primack, J. R., Dekel, A., et al. 2018, [ApJ](#), **858**, 114
- Hunter, J. D. 2007, [Comput. Sci. Eng.](#), **9**, 90
- Ibata, R. A., Gilmore, G., & Irwin, M. J. 1994, [Nature](#), **370**, 194
- Jiang, I.-G., & Binney, J. 1999, [MNRAS](#), **303**, L7
- Jones, E., Oliphant, T., Peterson, P., et al. 2001, [SciPy: Open source scientific tools for Python](#) [Online; Accessed 2019-04-19]
- Kelley, T., Bullock, J. S., Garrison-Kimmel, S., et al. 2019, [MNRAS](#), **487**, 4409
- Kingma, D. P., & Ba, J. 2015, [3rd International Conference on Learning Representations, ICLR 2015](#), San Diego, CA, USA, May 7–9, 2015, Conference Track Proceedings
- Kollmeier, J. A., Zasowski, G., Rix, H. W., et al. 2017, ArXiv e-prints [arXiv:1711.03234]
- Koppelman, H., Helmi, A., & Veljanoski, J. 2018, [ApJ](#), **860**, L11
- Kroupa, P. 2001, [MNRAS](#), **322**, 231
- Krumholz, M. R., & Gnedin, N. Y. 2011, [ApJ](#), **729**, 36
- Kunder, A., Kordopatis, G., Steinmetz, M., et al. 2017, [AJ](#), **153**, 75
- Lancaster, L., Koposov, S. E., Belokurov, V., Evans, N. W., & Deason, A. J. 2019, [MNRAS](#), **486**, 378
- Larkoski, A. J., Moult, I., & Nachman, B. 2020, [Phys. Rept.](#), **841**, 1
- Leitherer, C., Schaerer, D., Goldader, J. D., et al. 1999, [ApJS](#), **123**, 3
- Lowing, B., Wang, W., Cooper, A., et al. 2015, [MNRAS](#), **446**, 2274
- Majewski, S. R., Skrutskie, M. F., Weinberg, M. D., & Ostheimer, J. C. 2003, [ApJ](#), **599**, 1082
- Malhan, K., & Ibata, R. A. 2018, [MNRAS](#), **477**, 4063
- Marinacci, F., Pakmor, R., & Springel, V. 2014, [MNRAS](#), **437**, 1750
- Mateu, C., Bruzual, G., Aguilar, L., et al. 2011, [MNRAS](#), **415**, 214
- McKinney, W. 2010, in [Proceedings of the 9th Python in Science Conference](#), ed. S. van der Walt, 51
- Myeong, G. C., Evans, N. W., Belokurov, V., Sanders, J. L., & Koposov, S. E. 2018a, [ApJ](#), **863**, L28
- Myeong, G. C., Evans, N. W., Belokurov, V., Amorisco, N. C., & Koposov, S. E. 2018b, [MNRAS](#), **475**, 1537
- Myeong, G. C., Vasiliev, E., Iorio, G., Evans, N. W., & Belokurov, V. 2019, [MNRAS](#), **488**, 1235
- Necib, L., Ostdiek, B., Lisanti, M., et al. 2019a, ArXiv e-prints [arXiv:1907.07681]
- Necib, L., Ostdiek, B., Lisanti, M., et al. 2019b, ArXiv e-prints [arXiv:1907.07190]
- Necib, L., Lisanti, M., Garrison-Kimmel, S., et al. 2019c, [ApJ](#), **883**, 27
- Nissen, P. E., & Schuster, W. J. 2010, [A&A](#), **511**, L10
- Ofiorbe, J., Garrison-Kimmel, S., Maller, A. H., et al. 2014, [MNRAS](#), **437**, 1894
- Ostdiek, B., Necib, L., Cohen, T., et al. 2019, <https://doi.org/10.5281/zenodo.3354470>
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, [J. Mach. Learn. Res.](#), **12**, 2825
- Posti, L., Helmi, A., Veljanoski, J., & Breddels, M. A. 2018, [A&A](#), **615**, A70
- Robin, A., & Creze, M. 1986, [A&A](#), **157**, 71
- Sanderson, R. E., Garrison-Kimmel, S., Wetzel, A., et al. 2018, [ApJ](#), **869**, 12
- Sanderson, R. E., Wetzel, A., & Loebman, S. 2020, [ApJS](#), **246**, 6
- Searle, L., & Zinn, R. 1978, [ApJ](#), **225**, 357
- Sharma, S., Bland-Hawthorn, J., Johnston, K. V., & Binney, J. 2011, [ApJ](#), **730**, 3
- Springel, V. 2005, [MNRAS](#), **364**, 1105
- van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, [Comput. Sci. Eng.](#), **13**, 22
- Veljanoski, J., Helmi, A., Breddels, M., & Posti, L. 2019, [A&A](#), **621**, A13
- Wang, L., Dutton, A. A., Stinson, G. S., et al. 2015, [MNRAS](#), **454**, 83
- Weinberg, M. D. 1998, [MNRAS](#), **299**, 499
- Wetzel, A. R., Hopkins, P. F., Kim, J.-H., et al. 2016, [ApJ](#), **827**, L23
- White, S. D. M., & Rees, M. J. 1978, [MNRAS](#), **183**, 341
- Yosinski, J., Clune, J., Bengio, Y., & Lipson, H. 2014, [Advances in Neural Information Processing Systems](#), eds. Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, & K. Q. Weinberger (Curran Associates, Inc.), 27, 3320
- Zhang, Y., Mesaros, A., Fujita, K., et al. 2019, [Nature](#), **570**, 484

Appendix A: in situ stars

The distributions of the stars marked as in situ by the network are shown in Fig. A.1. The disk is easy to see in the l - b plane.

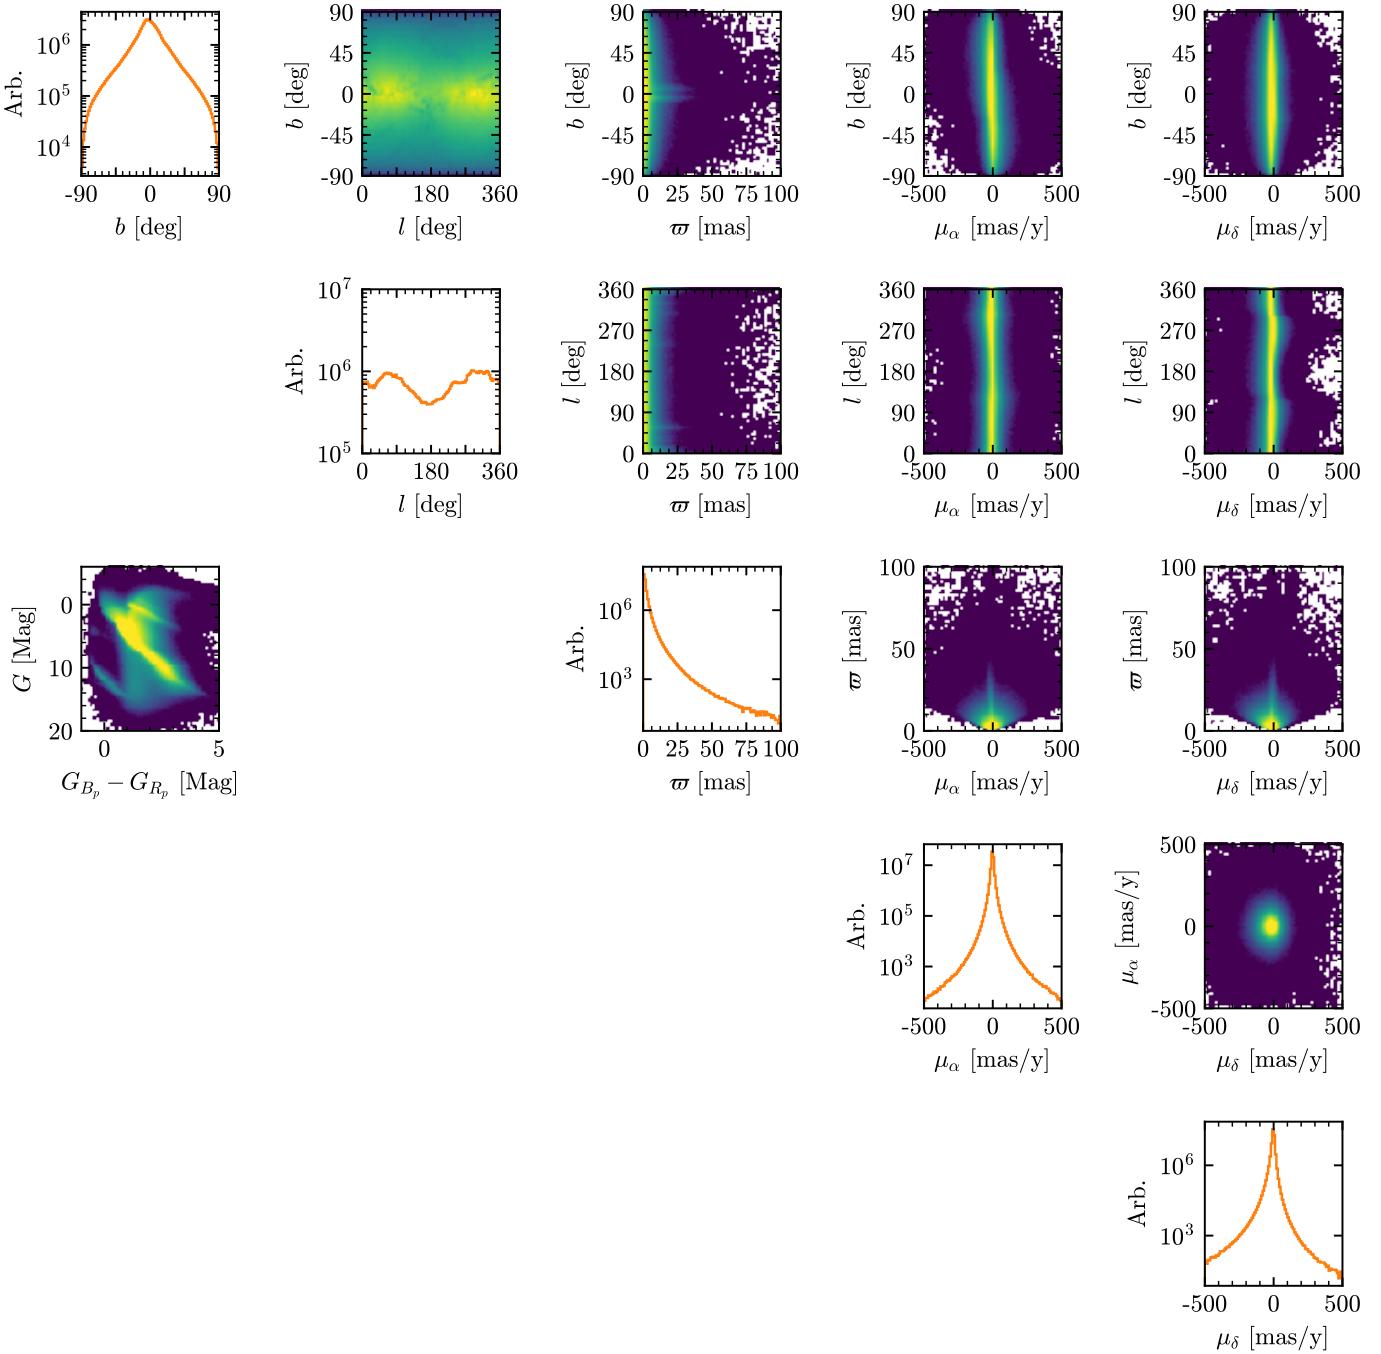


Fig. A.1. Stars within *Gaia* DR2 with $\delta\sigma/\sigma < 0.10$. These stars have been identified as in situ by our neural network. The b distribution is strongly peaked at 0 and the l distribution is relatively flat, consistent with stars that make up the disk. See Fig. 13 in the main text for the analogous distributions of the stars classified as accreted.

Appendix B: Accounting for measurement uncertainties

Given the finite resolution of *Gaia*, it is critical to incorporate the measurement uncertainties into our approach. Fortunately, the FIRE mock catalogs include uncertainties in the “measured” phase space of the simulated stars. This provides us with the ability to test to what extent it is important to include these uncertainties when training the neural network. In Fig. B.1, three stars from the test set of m12i LSR0 are shown. The distributions are made by randomly sampling the stars over their uncertainties and then applying the given network 200 times. The pairs of distributions shown for each star correspond to the neural network outputs for two different approaches to training. The solid lines are computed using the network that was trained without any error sampling in the training step. In the left and middle panels, the network is very uncertain, and we see there is support across the entire range between 0 and 1. In the right panel, there is still support over a large range, although most of the random samples fall above a score of 0.5.

The dashed lines show the networks with the configuration used in this work, in which error sampling is included as part of

the training step. Specifically, for each training epoch, a star in the training set is randomly sampled 20 times within its uncertainties. As the networks typically take around 50 epochs to finish training, the network ultimately sees around 1000 different random samplings per star. In the left panel, the star that the unsampled network was very uncertain about now has a strong preference for low network scores. In the middle panel, the network is still not very certain, but the star has lost support at scores as high as 1. In contrast to this, in the right panel, the star shifts to even larger network outputs after seeing many samples during training. We interpret this right panel as showing how a truth-level accreted star would be correctly classified by the network that was trained with error sampling, and would have been otherwise missed by the approach that did not include the impact of errors.

We note that the error sampling used here uses uncorrelated uncertainties, as provided in the FIRE catalogs. However, *Gaia* DR2 uncertainties are not so diagonal and include correlations. Given our tight parallax uncertainty cuts, we do not expect this to be an issue. Sampling uncorrelated errors during training should still lead to more precise network scores than not sampling.

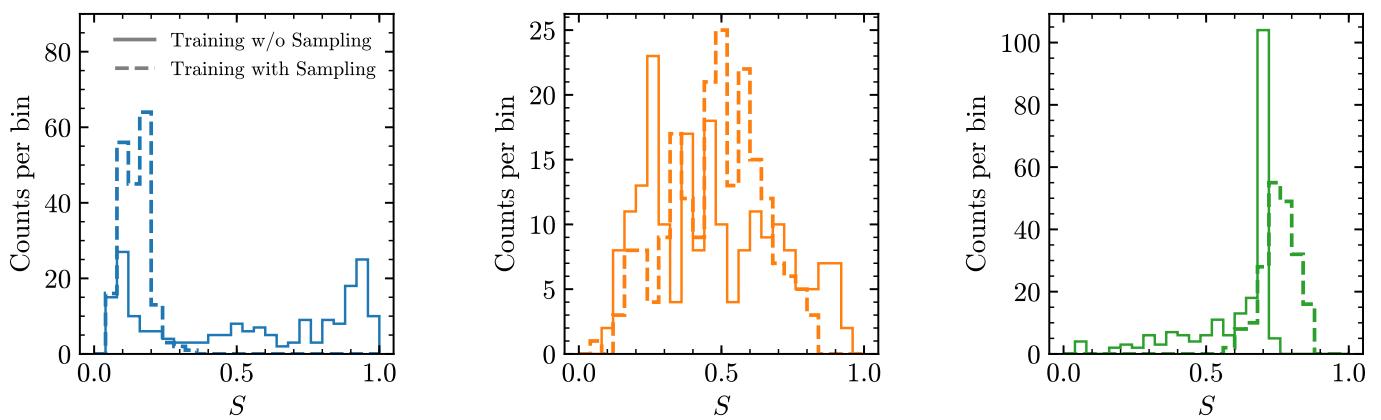


Fig. B.1. Each panel shows a star from the m12i LSR0 test set. Each of the three panels are computed by taking 200 samples of a single star, where the distribution is derived using the quoted uncertainties in the mock dataset. The dashed and solid lines show how the network output changes if it was trained with or without error sampling, respectively.

Appendix C: Kinematic + Photometric network results

In Sect. 5.2, we showed that the networks that had access to photometric information as an input were not able to generalize from bright, well measured stars (that is, those with 6D phase-space measurements) to dimmer stars without line-of-sight velocity measurements (that is, those with 5D phase-space measurements). This implied that only training on real stars for which we could infer empirical labels (implying that they are nearby and bright) would not work for extending to a more general catalog. However, with transfer learning, we are able to train on simulated stars that do not have a measurement of the line-of-sight velocity. This motivates constructing a catalog from a neural network that includes kinematic and photometric inputs.

Figure C.1 shows the results of training and testing kinematic + photometric networks on different viewpoints of the m12i galaxy, always using $\delta\varpi/\varpi < 0.10$, such that the photometric distributions should be similar. Unlike the kinematic networks (Fig. 6), the networks trained on multiple LSRs seem to always do better than training on either of the “wrong” LSRs alone. The fractional difference between the orange and purple lines is 8% and 12% at a false positive rate of 0.01, which is very similar to that of the kinematic networks. We conclude that although the photometric networks do not generalize well when extrapolating from brighter to dimmer stars, they do generalize as well as the kinematic-only networks when applied to different viewpoints.

Encouraged by these findings, we repeat the transfer learning methodology tests presented in Sect. 7 using 200 000 stars from m12f LSR0 that satisfy $\delta\varpi/\varpi < 0.10$ and have radial velocity measurements. The ZM selection is again found to be the best empirical label for transferring the pre-trained networks while only updating the last layer of the network. The optimal cut value to match the truth-level accreted star velocity distributions is 0.9. When the method is validated against m12f LSR1, the results are consistent.

For the final classifier that will be applied to the real *Gaia* data, we start with the kinematic + photometric networks

pre-trained on all of the stars from m12i LSR0, LSR1, and LSR2 with $\delta\varpi/\varpi < 0.10$. Transfer learning is performed using the stars of the RAVE DR5-*Gaia* DR2 cross-matched catalog, using the optimal method discussed above to determine empirical labels. Overall, the network using photometric information identifies about half as many accreted candidate stars as does the network only using kinematics, 383 k compared to 767 k. In addition, the photometric network seems to have more contamination from disk stars than expected. For example, Fig. C.2 shows the subset of stars with 3D velocity measurements. The most concerning feature is that each of the distributions in the top panels has a peak consistent with the disk dynamics. In the middle panels, where the 2D velocity distributions of the selected stars are shown, we observe large densities in the same location as the peaks in the in situ stars. This is reminiscent of what is observed in Veljanoski et al. (2019); although their boosted decision tree did not select these specific stars, they did tag a collection of stars whose colors are consistent with metal-poor stars but have kinematics consistent with the disk. We interpret these peaks as telling us that the kinematic + photometric network is relying too heavily on the photometric information, which could be pointing to an inconsistency between the mock catalogs and the real Milky Way stars tagged as accreted by this network.

Finally, distributions for all of the stars within the *Gaia* DR2 dataset that have $\delta\varpi/\varpi < 0.10$ and are classified by the kinematic + photometric network as accreted are shown in Fig. C.3. No coordinate transformations have been done, that is, these are the distributions of the inputs that the network is provided with (other than to the HR diagram). Comparing with the similar results for the kinematic network shown in Fig. 13, the kinematic + photometric network selects many more stars at larger parallax and small proper motions. The distribution in the HR diagram is also more restricted, and in the $l-b$ plane, there is certainly what looks to be a large disk-like component. We conclude that while there might be interesting features worth exploring within the catalog computed by the kinematic + photometric network, great care should be taken when interpreting any results inferred from it.

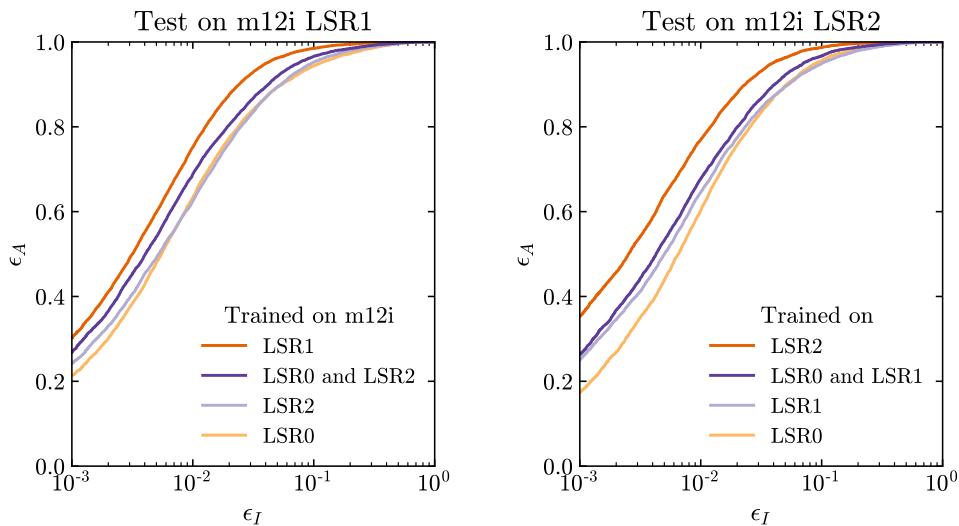


Fig. C.1. ROC curves showing the performance of networks trained on one LSR of m12i and applied to a different LSR. The networks are trained and tested on the subset of stars in the catalogs that have small parallax errors and use the photometric information in addition to the 5D kinematic information. The photometric data generalizes better between catalogs, so the reduction in performance is not as bad as the kinematics-only networks shown in Fig. 6.

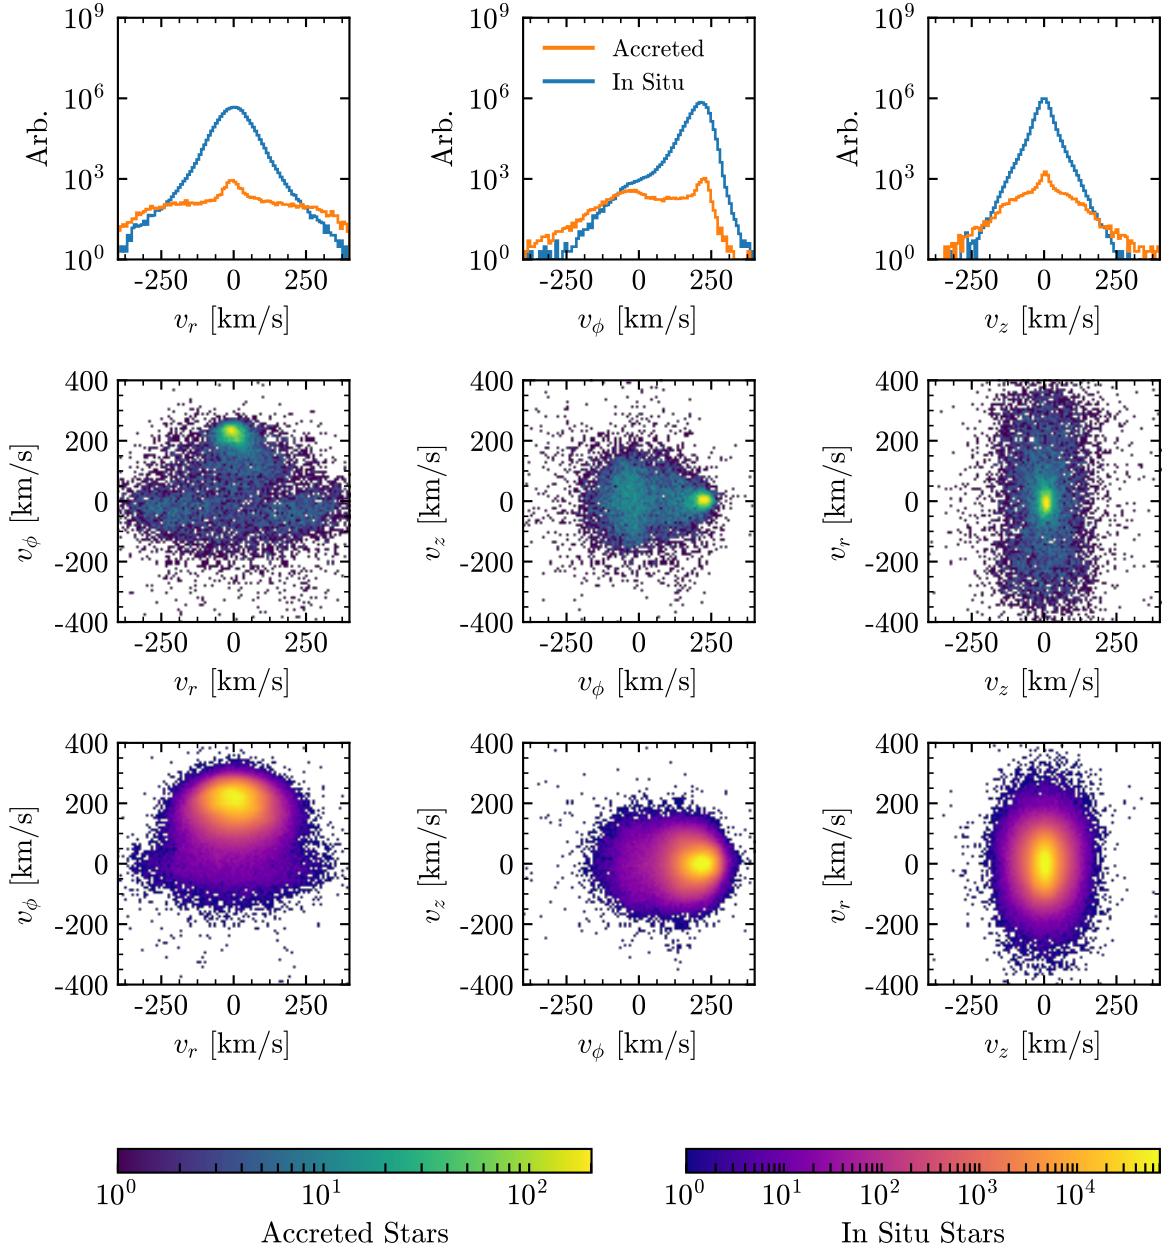


Fig. C.2. Galactocentric velocity distributions of the two sets of stars in the *Gaia* DR2 dataset with small parallax errors and a measurement of the radial velocity as classified by the kinematic + photometric network. The 2nd and 3rd rows provide correlation plots for the stars labeled as accreted and in situ by the kinematic + photometric network, respectively. Many of the selected stars have kinematics similar to the disk.

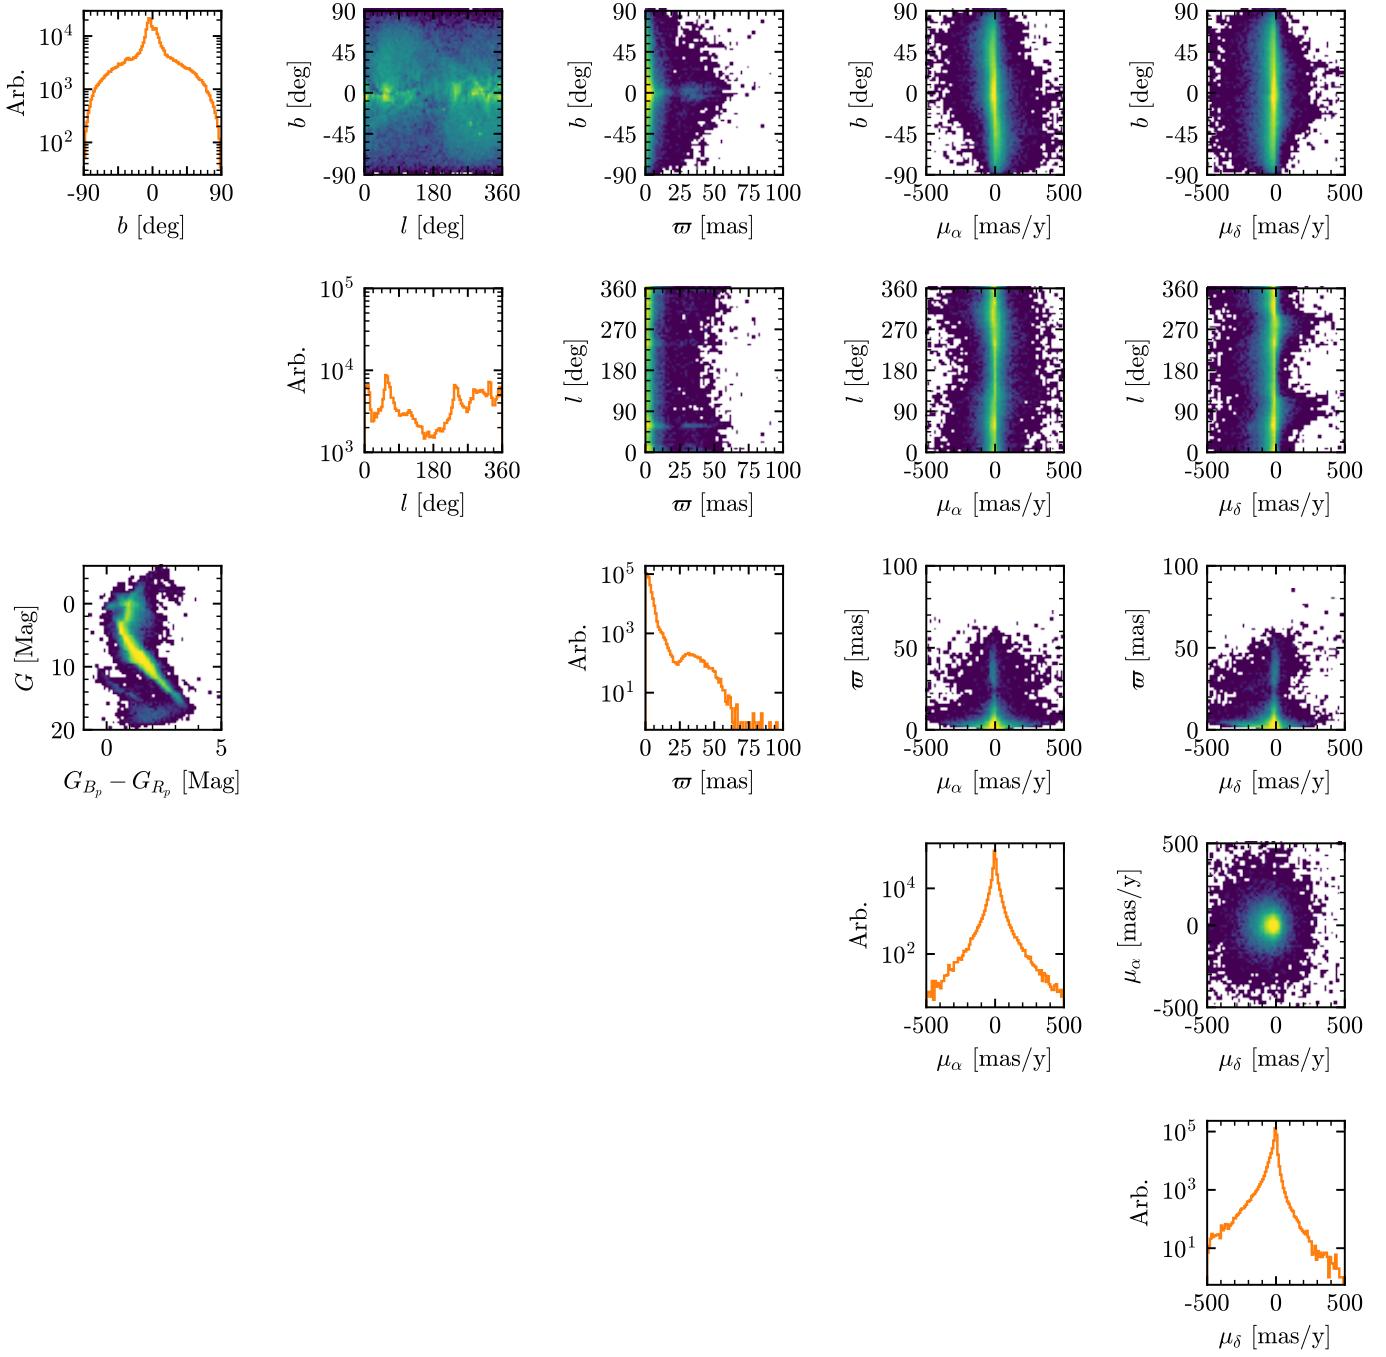


Fig. C.3. Stars from the entire set of stars within *Gaia* DR2 with $\delta\varpi/\varpi < 0.10$ that are classified as accreted by the kinematic + photometric network. See Fig. 13 for the analogous figure for the catalog derived using the kinematic network.