



Chasing Accreted Structures within Gaia DR2 Using Deep Learning

Lina Necib^{1,2,3} , Bryan Ostdiek⁴ , Mariangela Lisanti⁵, Timothy Cohen⁴, Marat Freytsis^{6,7}, and Shea Garrison-Kimmel⁸ 

¹Walter Burke Institute for Theoretical Physics, California Institute of Technology, Pasadena, CA 91125, USA; lincib@caltech.edu

²Department of Physics and Astronomy, University of California Irvine, Irvine, CA 92697, USA

³Observatories of the Carnegie Institution for Science, 813 Santa Barbara St., Pasadena, CA 91101, USA

⁴Institute of Theoretical Science, Department of Physics, University of Oregon, Eugene, OR 97403, USA; bostdiek@g.harvard.edu

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

⁶Raymond and Beverly Sackler School of Physics and Astronomy, Tel Aviv University, Tel-Aviv 69978, Israel

⁷School of Natural Sciences, Institute for Advanced Study, Princeton, NJ 08540, USA

⁸TAPIR, California Institute of Technology, Pasadena, CA 91125, USA

Received 2020 April 23; revised 2020 August 28; accepted 2020 September 12; published 2020 October 29

Abstract

In previous work, we developed a deep neural network classifier that only relies on phase-space information to obtain a catalog of accreted stars based on the second data release of Gaia (DR2). In this paper, we apply two clustering algorithms to identify velocity substructure within this catalog. We focus on the subset of stars with line-of-sight velocity measurements that fall in the range of Galactocentric radii $r \in [6.5, 9.5]$ kpc and vertical distances $|z| < 3$ kpc. Known structures such as Gaia Enceladus and the Helmi stream are identified. The largest previously unknown structure, Nyx, is a vast stream consisting of at least 200 stars in the region of interest. This study displays the power of the machine-learning approach by not only successfully identifying known features but also discovering new kinematic structures that may shed light on the merger history of the Milky Way.

Unified Astronomy Thesaurus concepts: Milky Way dynamics (1051); Galaxy dynamics (591); Astrometry (80); Neural networks (1933); Star clusters (1567)

1. Introduction

The paradigm of hierarchical structure formation describes how galaxies grow in a Lambda Cold Dark Matter universe (White & Rees 1978). During this process, large galaxies like our Milky Way gain the majority of their mass by capturing and absorbing smaller satellites. As a satellite galaxy falls onto a host, it is torn apart by violent tidal forces that strip both its dark matter halo and stars. Debris from this process is left scattered about the host galaxy as a fossil remnant of the accretion events. Assuming that not enough time has passed for these objects to fully virialize, such remnants manifest as distinctive features in phase space: stellar clumps, streams, and clouds of tidal debris (Johnston et al. 1996; Johnston 1998; Bullock & Johnston 2005; Robertson et al. 2005; Font et al. 2006, 2011). If such signatures can be identified, one can use their properties to reconstruct aspects of the host galaxy’s evolution.

The study of spatial and kinematic substructure in the Milky Way’s stellar halo has a long history. Various works have focused on modeling its large-scale components, discovering that the stellar halo can be separated into (at least) an inner and an outer component based on kinematics based on distinctions in spatial density profiles, stellar orbits, and spectroscopic metallicities (Carollo et al. 2007, 2010; Beers et al. 2012), photometry (Hess diagrams) in combination with spatial density distributions (de Jong et al. 2010), and photometric metallicity estimates in combination with kinematics inferred from proper motions alone (An et al. 2013, 2015). These works targeted the overall structure of the stellar halo. In parallel, other studies focused on identifying the substructures that built the accreted halo. The Sagittarius Stream provides the most stunning example of an ongoing merger (Johnston et al. 1995; Ivezić et al. 2000; Yanny et al. 2000; Ibata et al. 2001), as it traces several orbits of the infalling Sagittarius dwarf (Ibata et al. 1994). Numerous tidal streams, some of which may be

associated with disrupting globular clusters, have also been discovered; see Newberg & Carlin (2016) for a review. These include the GD-1 (Grillmair & Dionatos 2006), Pal-5 (Odenkirchen et al. 2001), and Ophiuchus (Bernard et al. 2014) streams. In the case of older mergers, velocity coherence is still preserved although spatial coherence may be lost (Helmi & White 1999; Lisanti & Spergel 2012; Lisanti et al. 2015). A plethora of such kinematic substructure has been found throughout the stellar halo, e.g., ECHOS (Schlaufman et al. 2009, 2011, 2012).

The advent of data from the Gaia satellite (Brown et al. 2018) has already revolutionized our ability to reconstruct the Milky Way’s history, given the unprecedented number of 5D and even 6D stellar phase-space measurements. To date, several additional stellar streams and clumps have been discovered using Gaia data (Koppelman et al. 2018, 2019; Malhan & Ibata 2018; Myeong et al. 2018a, 2019; Ibata et al. 2019a, 2019b; Meingast et al. 2019). Gaia Enceladus (also referred to as the Gaia Sausage) is the largest of these new structures (Belokurov et al. 2018; Helmi et al. 2018) and is believed to be the remnant of a significant merger that occurred at an estimated redshift of $\sim 1-3$ (Myeong et al. 2018b). Today, Enceladus’ stellar debris is highly radial and more metal rich than the rest of the stellar halo; additionally, it constitutes the majority of the accreted stellar fraction in the inner Milky Way (Deason et al. 2018; Lancaster et al. 2019; Necib et al. 2019a; Vincenzo et al. 2019).

The focus of this paper is on searching for kinematic substructures in the subset of Gaia data that includes radial velocities. In particular, we define the region of interest (ROI) to be within spherical Galactocentric radii $r \in [6.5, 9.5]$ kpc and vertical distances $|z| < 3$ kpc. In this region, nearly 99% of all stars belong to the Milky Way’s disk, while the remaining $\lesssim 1\%$ belong to the halo and may have been accreted from mergers (Jurić et al. 2008). Identifying this small fraction of accreted stars is a daunting task given the considerable size of

the stellar disk background population. A variety of strategies have been explored to achieve this goal, which include looking for structures in action-angle (Yuksel et al. 2009), energy–momentum (Helmi & de Zeeuw 2000), apocenter–pericenter–angular-momentum (Helmi et al. 2006), and/or chemodynamic (e.g., Nissen & Schuster 2010; Helmi et al. 2017; Posti et al. 2018; An & Beers 2020) space.

The Helmi stream provided the first proof of principle that such techniques can be applied successfully (Helmi et al. 1999). First discovered as an overdensity of 13 stars in angular-momentum space, the Helmi stream has since been confirmed by a variety of other observations (Chiba & Beers 2000; Re Fiorentin et al. 2005; Klement et al. 2009; Smith et al. 2009; Beers et al. 2017). Most recently, it has been studied in Gaia DR2 (Gaia Collaboration et al. 2018b; Koppelman et al. 2018) and nearly 600 more potential members were identified using cross-matches with spectroscopic surveys (Koppelman et al. 2019). Comparisons with simulations strongly suggest that it is associated with the disruption of a satellite galaxy (Kepley et al. 2007).

The selection criteria typically used to identify accreted stars are conservative and usually reduce potential disk contamination at the expense of excluding a large fraction of stars. Ideally, one would want to refine the selection process to yield a high-purity sample of accreted stars without sacrificing the overall statistics. Attempts in this direction have been made using traditional regression and classification techniques (Malhan & Ibata 2018; Veljanoski et al. 2019). Our focus here is on the application of deep neural networks to this problem. In Ostdiek et al. (2020), we developed a scheme for using neural networks to successfully distinguish accreted stars from those born in the Milky Way. We validated this approach through extensive testing on simulated Gaia mock catalogs (Sanderson et al. 2020) based on the Latte suite (Wetzell et al. 2016) of the FIRE simulations (Hopkins 2015; Hopkins et al. 2018). Our method is applied to the subset of Gaia stars with small parallax errors $\delta\varpi/\varpi < 0.10$ and yields a catalog of likely accreted stars (publicly available in Ostdiek et al. 2019). This work presents the first analysis of a subset of this catalog, focusing on the 4.8 million stars that additionally have radial velocity measurements.

We use several different clustering algorithms to identify velocity structures in this new catalog. We recover known structures such as the Helmi stream and Gaia Enceladus, clearly showing that the latter extends down to the Galactic midplane. In addition, we identify several new stream candidates in the ROI. One of these candidates, which we call Nyx, is a significant contributor in the region studied (Necib et al. 2020). The other two stream candidates, which may be associated with overdensities in Koppelman et al. (2018), comprise $\mathcal{O}(10)$ stars.

This paper is organized as follows. In Section 2, we review the methodology underlying the catalog obtained by Ostdiek et al. (2020) and explicitly demonstrate how neural network scores correlate with different stellar populations (e.g., thin/thick disk and halo). In Section 3, we perform a Gaussian mixture model analysis to characterize the three most significant structures in the sample: Gaia Enceladus, Nyx, and the remaining constituents of the stellar halo, which for the remainder of this paper we will refer to as “Halo.” We further analyze the data for non-Gaussian structures in Section 4, which is where we recover the Helmi stream, as well as two

other stream candidates. The Appendix includes additional figures that validate the analysis procedure.

2. A New Catalog of Accreted Stars

In this section, we take a first look at the catalog of accreted stars. After briefly summarizing the methodology that underlies the catalog, we provide a number of distributions of the basic properties of the stars identified as accreted and in situ by the neural network. Intriguingly, one can already see hints of novel structures by looking at these distributions. We will explore these structures more quantitatively in Section 3.

2.1. Characterizing the Catalog

To take maximum advantage of the richness contained within Gaia DR2, we developed a novel approach utilizing deep neural networks to derive a catalog of accreted stars. Details and extensive cross-checks of the methodology are presented in Ostdiek et al. (2020); we briefly summarize the main points here. We train the neural network on a combination of simulated mock Gaia surveys and real measurements of Milky Way stars. The simulated data is from the Ananke mock surveys (Sanderson et al. 2020), based on the Latte simulation suite (Wetzell et al. 2016; Hopkins et al. 2018). As we have the full merger history for each simulated galaxy, we can identify stars that are truly accreted, which can be leveraged to both train and validate the networks. The final network is pretrained on mock catalogs from three solar positions of the `m12i` simulated galaxy, using only 5D heliocentric kinematics as inputs (the line-of-sight velocity is not provided to the networks). We then perform transfer learning on stars within the RAVE DR5–Gaia DR2 cross-matched data (Kunder et al. 2017) that are identified as accreted with high confidence, i.e., if they are metal poor $[\text{Fe}/\text{H}] < -1.5$ and have $|z| > 1.5$ kpc. Retraining the last layer of the neural network allows it to learn characteristics specific to the Milky Way.

The final network is applied to the subset of Gaia DR2 stars with parallax $\varpi > 0$ and fractional error $\delta\varpi/\varpi < 0.10$. In this work, we further restrict our analysis to only include those stars in the resulting catalog that have measured Gaia line-of-sight velocities (Katz et al. 2019) and fall within spherical Galactocentric radii of $r \in [6.5, 9.5]$ kpc and vertical distances $|z| < 3$ kpc of the Galactic midplane. This results in a final sample size of 4,820,164 stars. Figure 1 shows the spatial extent of the sample analyzed here. Note that Gaia DR2 is essentially complete in the optical magnitude range $G \in [7, 17]$ (Gaia Collaboration et al. 2018a). For $G \in [4, 12]$, the radial velocity subset of Gaia is about 60%–80% complete as compared to the total data set. Therefore, we do not expect our sample to be spatially complete, and this is evident from close inspection of Figure 1. For example, the distribution peaks at the solar circle, and there are rays originating from the solar position $(x, y, z) = (-8, 0, 0)$ kpc in the x – y plane. There is also an asymmetry in the y – z plane where the distribution looks almost diagonal, probably due to projection effects (see Katz et al. 2019 for the completeness of the radial velocity subset of Gaia DR2). Although the stellar sample is not spatially complete, we do expect it to be kinematically unbiased because no selection cuts have been made on velocity. This is particularly important for the study presented here, as the techniques we use for identifying stellar substructure are based on features in velocity space alone.

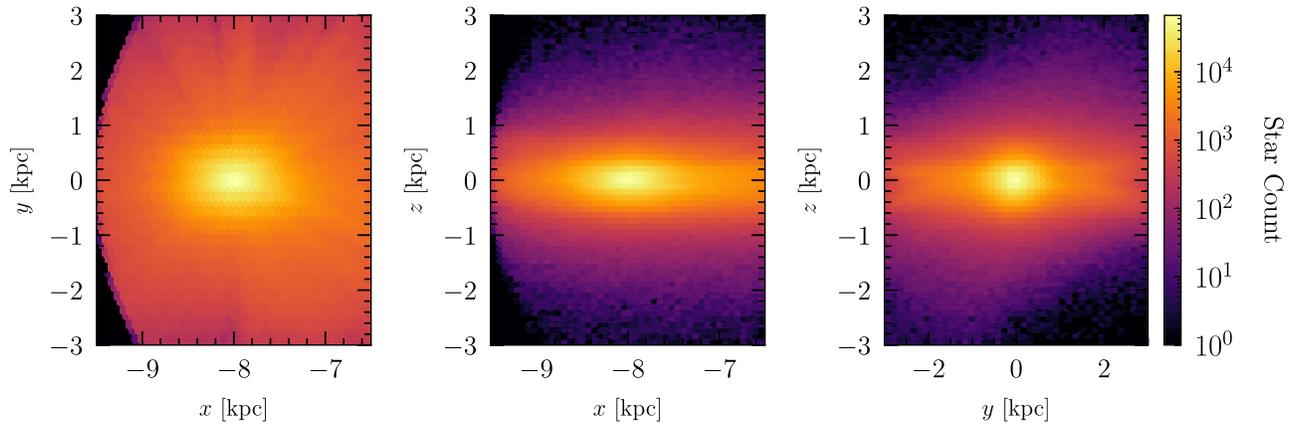


Figure 1. Spatial distribution of all Gaia DR2 stars with parallax $\varpi > 0$, fractional error $\delta\varpi/\varpi < 0.10$, and line-of-sight velocities, falling within the region $r \in [6.5, 9.5]$ kpc and $|z| < 3$ kpc. Coordinates are in the Galactocentric Cartesian frame, with $z = 0$ the Galactic midplane, $y > 0$ pointing in the direction of the Sun’s motion, and $x = 0$ at the Galactic Center. The Sun is located at $(x, y, z) = (-8, 0, 0)$ kpc.

The neural network gives each star a score that reflects its probability of being accreted, with $S = 0$ indicating a maximum-confidence in situ star and $S = 1$ a maximum-confidence accreted star. Figure 2 shows the Toomre plot in Galactocentric coordinates of the stars in Gaia DR2 that pass all the selection cuts discussed above. A gradient of scores is evident, ranging from stars associated with the thin disk ($S \sim 0$) to stars associated with the stellar halo ($S \sim 1$). For context, we additionally overlay the best-fit 3σ contours for the thin disk, thick disk, and halo stars derived from the model of Bensby et al. (2003).

Figure 2 clearly shows that the network scores are highly correlated with the expected behavior of these three populations: the lowest scores are associated with the thin disk, while the larger scores are associated with halo stars. For the most part, thin-disk stars have network scores $S \lesssim 0.05$, while thick-disk stars have scores extending up to $S \sim 0.5^9$ —see Figure A1 for Toomre plots of stars with scores $S < 0.05$ and $S \in [0.3, 0.5]$. For values of $v_y > 200$ km s $^{-1}$ and $\sqrt{v_x^2 + v_z^2} \sim 100$ – 200 km s $^{-1}$, stars that appear to belong to the thick disk according to the Bensby et al. (2003) model have very low network scores, close to 0. These stars are easily identified by the network as being in situ because their v_y velocities are large and positive. Note that due to the focus of the training regimen, the network score reflects the probability that a star is accreted, and as such, it does not have any information on the different components of the disk.¹⁰ It is therefore rather remarkable that the network scores do correlate with known stellar populations.

To create a concrete data set of accreted stars to work with, we must choose a cut on the network scores. In Ostdiek et al. (2020), we developed a principled way of identifying this cut by analyzing the network performance on mock catalogs, where truth information is available. Our criterion was to minimize the difference between velocity distributions of the stars that are selected and the distributions of all of the

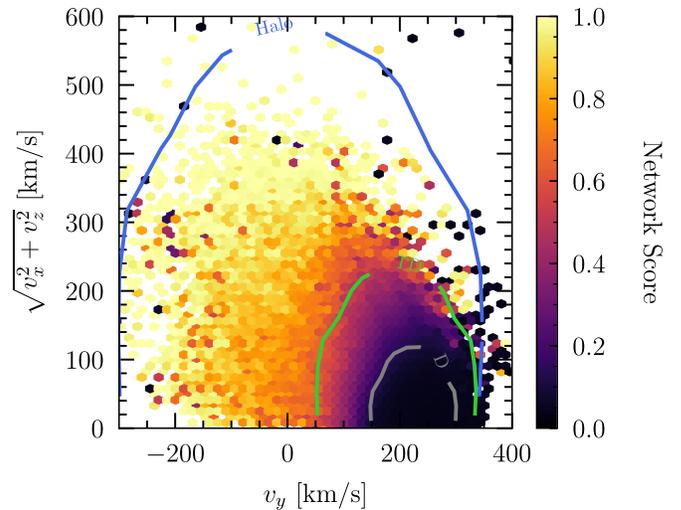


Figure 2. Toomre plot of all stars in the subset of Gaia DR2 with line-of-sight velocities, $\delta\varpi/\varpi < 0.10$, and $\varpi > 0$ in the region defined by $r \in [6.5, 9.5]$ kpc and $|z| < 3$ kpc. Here, positive v_y points in the same direction as the disk rotation. The colors of the points indicate their network score. Overlaid are the 3σ velocity contours of the thin disk (D), thick disk (TD), and halo stars, obtained using the model of Bensby et al. (2003). The network scores are clearly correlated with the different stellar populations.

truth-level accreted stars. In particular, we found that a cut of $S \gtrsim 0.75$ least biased the distributions of the accreted stars.¹¹ Here, we will use more restrictive cuts on the score to improve the purity of the sample and facilitate the Gaussian Mixture Model analysis described in Section 3.1. Our canonical sample will be defined with $S > 0.85$; it contains 22,296 stars total. We additionally explore the impact of placing a much more restrictive cut of $S > 0.95$ to create a high-purity sample (9379 stars total). The high-purity sample minimizes disk contamination at the expense of biasing the velocity distributions, while the canonical sample reduces the purity by introducing some disk stars but should produce less biased distributions.

⁹ Running the neural network on the *Galaxia* (Sharma et al. 2011) simulation demonstrated that the thick-disk component is the hardest one to score as accreted or in situ. In actuality, the Milky Way’s thick disk might be a combination of the two; see Rix & Bovy (2013) and references therein.

¹⁰ It is in principle possible to use multi-class classification to identify different stellar populations, but that is outside the scope of this work.

¹¹ All stars in Gaia with a score greater than 0.75 are made available upon request from the corresponding authors.

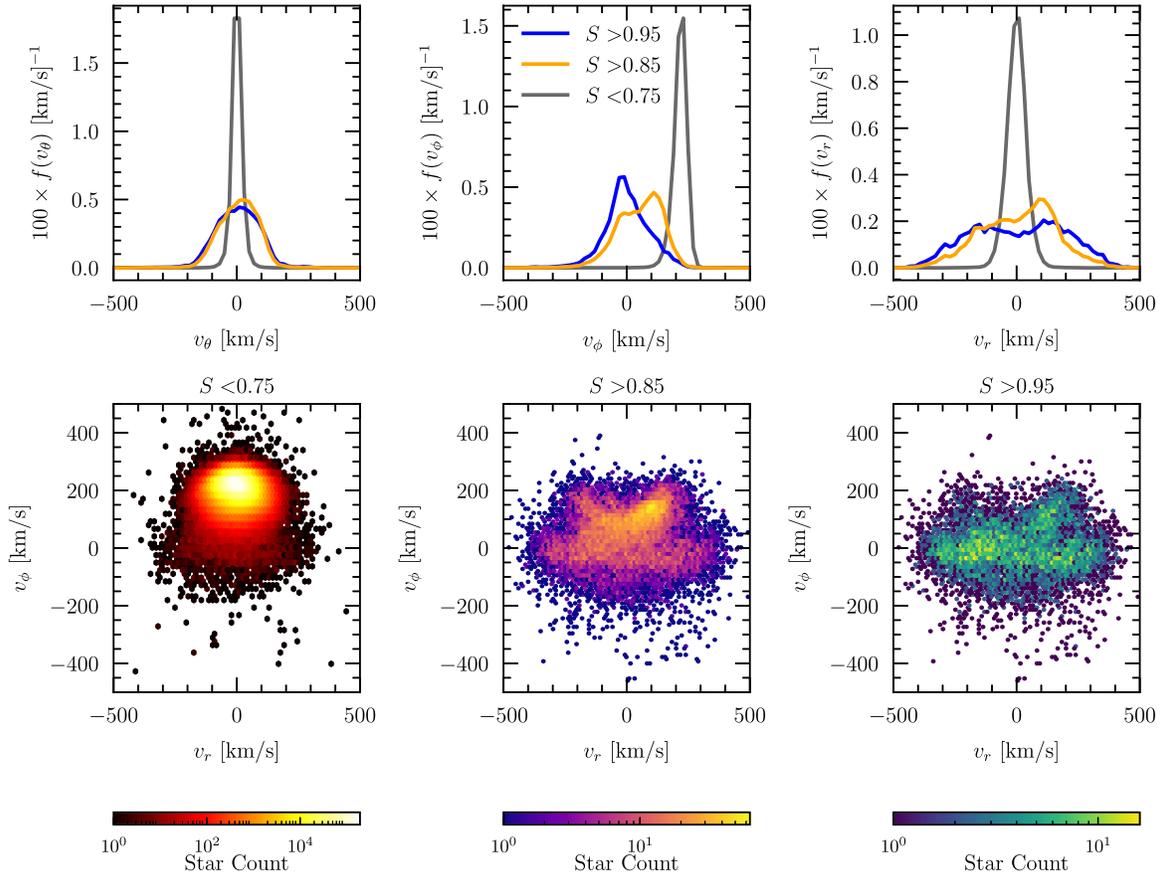


Figure 3. The top row shows the spherical Galactocentric velocity distributions of the stars in the Gaia DR2 data set that pass the cuts described in Section 2.1. Here, v_ϕ rotates with the disk of the Milky Way and positive v_r points toward the Galactic Center. Distributions are provided for the in situ ($S < 0.75$), canonical ($S > 0.85$), and high-purity ($S > 0.95$) samples. The second row provides $v_r - v_\phi$ correlation plots for these three samples. The in situ sample is clearly peaked at the disk rotation ($v_\phi \sim 220 \text{ km s}^{-1}$). Gaia Enceladus is apparent in the canonical and high-purity samples as the radial overdensity at $v_\phi \sim 0 \text{ km s}^{-1}$. As we will demonstrate, the stellar overdensity at $v_r \sim 130 \text{ km s}^{-1}$ and $v_\phi \sim 130 \text{ km s}^{-1}$, which is most apparent in the high-purity sample (but also present in the canonical sample), corresponds to a new stream, which we call Nyx.

2.2. Accreted versus In Situ Stars

We now explore the kinematic distributions of the canonical and high-purity samples. The top row of Figure 3 shows the velocity distributions in spherical Galactocentric coordinates (v_r , v_ϕ , v_θ) for three different cuts on the network score: $S < 0.75$ (in situ), $S > 0.85$ (canonical accreted), and $S > 0.95$ (high-purity accreted). The narrow peak near $v_\phi \sim 200 \text{ km s}^{-1}$ for the stars with $S < 0.75$ shows that the network clearly classifies stars whose rotations are consistent with the disk as in situ. The v_r and v_ϕ distributions differ between the canonical and high-purity sample. While both are roughly bimodal in v_r , the canonical sample has a smaller dispersion. The canonical sample has a stronger peak at $v_\phi \sim 150 \text{ km s}^{-1}$, which may due to some disk contamination.

To further explore the differences between the canonical and the high-purity cuts, we provide the $v_r - v_\phi$ distributions of the resultant data sets in the bottom row of Figure 3. The in situ sample ($S < 0.75$) is dominated by the stars at $v_r \sim 0 \text{ km s}^{-1}$ and $v_\phi \sim 200 \text{ km s}^{-1}$ as expected, although it also includes some accreted stars that happen to be scored below $S = 0.75$ by the network. Both the high-purity and canonical samples show an elongated structure in the radial direction that spans $v_r \in [-400, 400] \text{ km s}^{-1}$, with $v_\phi \sim 0 \text{ km s}^{-1}$. This is Gaia Enceladus, further discussed in Section 3.2. When we

implement the canonical cut, we find that the network carves out a nearly circular region in this parameter space, because any star that falls within this region is very likely to be in situ. It does, however, leave a distinct half-moon structure in this plane, particularly apparent in the middle panel of Figure 3, second row. If we increase the cut to $S > 0.95$, so that we only plot the stars that the network labels as accreted with high confidence, we find that the half-moon structure dissolves into a localized overdensity at $v_r \sim 150 \text{ km s}^{-1}$ and $v_\phi \sim 140 \text{ km s}^{-1}$. We call this new structure Nyx, and will describe it in more detail in Section 3.2. It is important to note that Nyx is present in the canonical sample as well, but is much easier to identify in the high-purity sample.

Figure 4 provides the Toomre plots for the canonical and high-purity cuts. The black dashed lines are centered around $v_{\text{lsr}} = 220 \text{ km s}^{-1}$ and extend to $|v - v_{\text{lsr}}| > 100\text{--}220 \text{ km s}^{-1}$. As with the $v_r - v_\phi$ distributions shown in Figure 3, the network cuts out the region within the inner dashed circle. A common selection criterion for identifying accreted stars is to require that $|v - v_{\text{lsr}}| > 180\text{--}220 \text{ km s}^{-1}$ (see Posti et al. 2018 for a review). By default, such a cut ignores stars that fall within (approximately) the outermost dashed circle in Figure 4, biasing the resulting distribution of stars in velocity space. This consequently ignores any substructure that corotates with the disk. Our deep-learning-based catalog has been constructed to

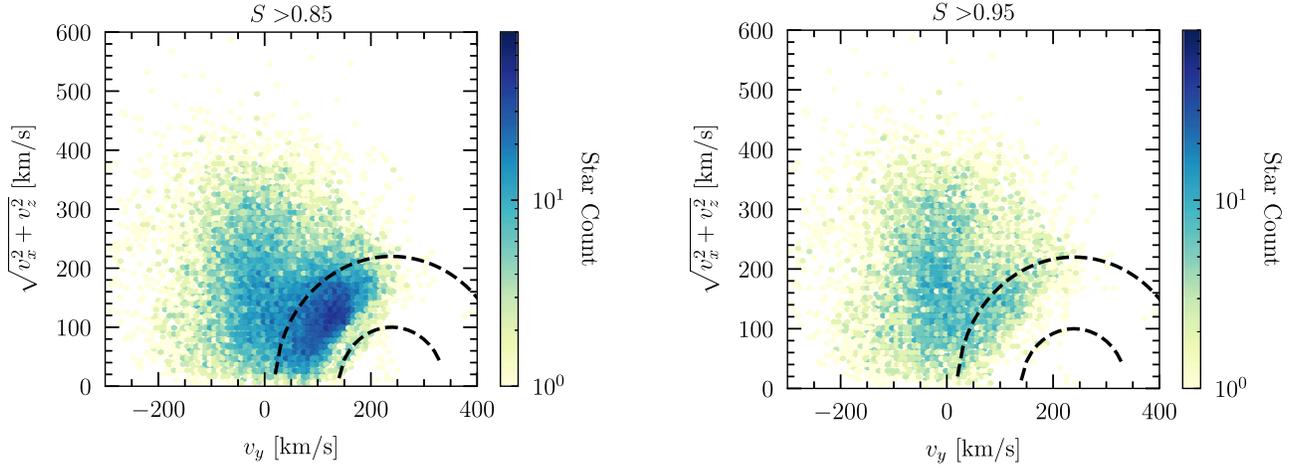


Figure 4. Toomre plots for the stars in the accreted catalog with the canonical (left) and high-purity (right) samples. The black, dashed lines show velocities within 100 and 220 km s^{-1} of the local frame. The network does not classify any stars that are within the inner dashed circle as accreted, and it has higher confidence that stars that are farther away from the local standard of rest are accreted.

minimize such a bias, thereby enabling us to look for prograde structures. Several structures stand out in both the canonical and high-purity samples. The vertically extended overdensity centered at $v_y \sim 0 \text{ km s}^{-1}$ is Gaia Enceladus. Second, the overdensity centered at $v_y \sim 130 \text{ km s}^{-1}$ and $\sqrt{v_x^2 + v_z^2} \sim 130 \text{ km s}^{-1}$ is Nyx. Nyx becomes more evident in the high-purity sample where the disk contamination is reduced.

We emphasize that the kinematic distribution of the stars that are easiest to label as accreted are not guaranteed to be representative of the distribution for *all* accreted stars. Therefore, moving forward, we will present results for both the canonical and high-purity samples, although we emphasize the possibility of biases in the distributions, especially as the cut on S becomes more restrictive (Ostdiek et al. 2020).

3. Dominant Kinematic Structures

In this section, we focus on characterizing the largest (in terms of overall star count) kinematic structures present in the sample of accreted stars. We will use a Gaussian mixture model approach to extract three distinct features, including the new stream Nyx.

3.1. A Structure-finding Algorithm

To obtain a quantitative estimate for the number of components in the high-purity accreted data set (defined by the sample of Section 2.1 with $S > 0.95$), we run a Gaussian mixture model using `scikit-learn` (Pedregosa et al. 2011) on the selected stars in Galactocentric velocity space. This first pass ignores the impact of finite measurement errors on the stellar velocities. We allow for up to nine Gaussian distributions and evaluate how many are necessary using the Bayesian Information Criterion (BIC). The BIC decreases rapidly as the number of Gaussians is increased to four, after which it stabilizes. Note that Gaia Enceladus is best characterized by two Gaussians with nearly identical means and dispersions in v_ϕ and v_θ that are centered at equal but opposite values of v_r (Necib et al. 2019a). Hence, the fact that four Gaussians are required to minimize the BIC tells us that the data set contains three dominant kinematic structures. We will refer to these three structures as Enceladus, Nyx, and the Halo. We

emphasize that these three components taken together comprise the totality of the accreted stellar halo in our model.¹² In other words, the component labeled as ‘‘Halo’’ is comprised of accreted stars in the stellar halo (some of which might be in coherent structures) that are not resolved by the mixture modeling.

The simple mixture analysis that yielded the BICs on the high-purity sample is efficient but does not account for the uncertainties in the measured parameters of the stars. The preliminary analysis informs the number of Gaussian distributions we include moving forward. To properly account for the errors, we perform our own dedicated study to derive the best-fit parameters of the velocity distributions of these separate components. We define the likelihood that a star i belongs to either the Halo or Nyx by

$$p_k(O_i | \theta) = \mathcal{N}(v_i | \mu^k, \Sigma_i^k), \quad (1)$$

where $k = h$ (Halo) or n (Nyx), $O_i = \{v_{r,i}, v_{\theta,i}, v_{\phi,i}\}$ are the velocities of star i in spherical Galactocentric coordinates, and θ is the set of free model parameters. \mathcal{N} denotes the multivariate normal distribution with mean $\mu = (\mu_r, \mu_\theta, \mu_\phi)$ and covariance Σ_{true} , which is a function of the dispersions $\sigma_{r,\theta,\phi}$ and the correlation coefficients $\rho_{r\theta,r\phi,\theta\phi}$. The measurement errors are taken into account as $\Sigma_i^k = \Sigma_{\text{true}}^k + \Sigma_{\text{err},i}$, where $\Sigma_{\text{err},i}$ varies from star to star but does not depend on the model parameters. The likelihood for Enceladus is the sum of two Gaussians with opposite means in the radial direction:

$$p_e(O_i | \theta) = \frac{1}{2} [\mathcal{N}(v_i | \mu^e, \Sigma_i^e) + \mathcal{N}(v_i | \mu^{\bar{e}}, \Sigma_i^{\bar{e}})], \quad (2)$$

where $\mu^{\bar{e}} = (-\mu_r, \mu_\theta, \mu_\phi)^e$. The total likelihood is therefore defined as

$$p(\{\mathcal{O}_i\} | \theta) = \prod_{i=1}^N \sum_{j=h,n,e} \mathcal{Q}_j p_j(O_i | \theta), \quad (3)$$

¹² There is possible disk contamination even in the high-purity sample. However, these stars are usually not representative of the in situ sample, and therefore, we do not expect them to form a distribution that can be picked up by the Gaussian mixture model.

Table 1

Parameters and Associated Prior Ranges for the Halo, Gaia Enceladus, Nyx, and Nyx-2

	Priors			
	Halo	Enceladus	Nyx	Nyx-2
μ_r	[-70, 70]	[0, 250]	[-70, 200]	[-300, 200]
μ_θ	[-70, 70]	[-200, 200]	[-70, 70]	[-70, 70]
μ_ϕ	[-70, 70]	[-200, 200]	[0, 300]	[0, 300]
$\sigma_{r,\theta,\phi}$	[0, 400]	[0, 400]	[0, 400]	[0, 400]
$\rho_{r\theta,r\phi,\theta\phi}$	[-1, 1]	[-1, 1]	[-1, 1]	[-1, 1]
Q	...	[0, 1]	[0, 1]	[0, 1]

Note. All priors are taken to be linear. Note that the Nyx-2 component is only included when analyzing the canonical ($S > 0.85$) sample.

where Q_j is the fractional contribution of each component (constrained to add up to 1 over all j). This type of likelihood analysis is very similar to what was performed in Necib et al. (2019a), except that here we only cluster in kinematic space and do not include metallicity information in the likelihood.

We also present results using the $S > 0.85$ sample. In this case, the total likelihood in Equation (3) includes an additional Gaussian, modeled following Equation (1), intended to capture a second component of Nyx at negative radial velocity, which we refer to as Nyx-2.¹³

We run *emcee* (Foreman-Mackey et al. 2013) to find the posterior distributions for the free parameters of the separate components. With 9 free parameters for each population and 2 additional parameters to quantify the relative fractions, this is a 29 parameter fit for the $S > 0.95$ sample study. With 4 components and 3 relative fractions, the $S > 0.85$ sample requires 39 free parameters. The priors are linear, and their ranges are listed in Table 1. We use 200 walkers, 8000 steps for the burn-in stage, and 1000 steps for the analysis of the high-purity (canonical) samples. The relevant corner plots are provided in the Appendix as Figures A2–A8.

3.2. Properties of Enceladus and Nyx

In this section, we investigate the properties of the dominant kinematic structures in the catalog. The top (bottom) row of Figure 5 provides the 2D kinematic distributions of the stars in the $S > 0.95$ (0.85) catalog in Galactocentric spherical coordinates. The lines show the 68% contours of the posterior distributions for the separate stellar components, modeled using the likelihood in Equation (3). The corresponding 1D distributions for the posteriors are shown in Figure 6.

Gaia Enceladus is the radial bimodal population in Figure 5 (blue line). By construction, the lobes of Enceladus have the same mean and dispersion in v_θ and v_ϕ and are located at equal, but opposite, values of the radial velocity; see Equation (2). The general properties of the Enceladus distribution are largely consistent with results from a previous study using the SDSS–Gaia cross-match (Necib et al. 2019a). From the radial distributions shown in Figure 6, we see that the peaks are located closer together in the canonical sample than in the high-purity sample. These differences may be a result of kinematic biases that are introduced as the score cut is increased.

¹³ Using the same technique of evaluating the BICs, we find that the number of preferred Gaussians stabilizes after six. We choose to include five Gaussians (two of which correspond to Gaia Enceladus). Adding a sixth Gaussian leads to two overlapping distributions in velocity space, roughly at the location of Nyx.

Additionally, as shown in Figure A10, we find that stars with a high probability of being associated with Enceladus clearly extend down to the Galactic midplane. This corroborates the hypothesis that the Enceladus merger contributes to the local dark matter distribution (Necib et al. 2019a, 2019b).

Nyx is the prograde group of stars characterized by its significant radial velocity (Figure 5, green line). In particular, Nyx moves in the same direction as the Galactic disk, but its rotational velocity lags by ~ 90 km s⁻¹. Its radial velocity distribution has a mean value of 134 km s⁻¹ and dispersion of ~ 67 km s⁻¹. In the canonical sample, we find a corresponding structure, which we call Nyx-2, that has an average v_ϕ of 94 km s⁻¹, compared to the Nyx value (in the same sample) of 125 km s⁻¹, but an equal and opposite v_r of -79 km s⁻¹ compared to 98 km s⁻¹ for Nyx.

The fraction of Nyx stars is 9% and 30% in the high-purity and canonical samples, respectively. Nyx-2 corresponds to 22% of the canonical sample. Enceladus is the dominant structure in both samples, comprising 77% (42%) of the high-purity (canonical) set. The relative fraction of Enceladus is reduced in the canonical sample primarily because of the presence of Nyx-2.

A complete discussion of the properties of Nyx and Nyx-2, and their potential origin, is presented in Necib et al. (2020). Briefly summarizing what is detailed there, Nyx is a coherent stellar stream on an eccentric orbit ($e \sim 0.6$) whose distribution is highly unlikely to be drawn from the expected smooth distribution of the thick disk. Its behavior is consistent with prograde streams observed in simulations, which are created when a massive satellite is dragged into the disk plane by increased tidal friction (Quinn & Goodman 1986; Lake 1989; Walker et al. 1996; Abadi et al. 2003; Read et al. 2008, 2009; Purcell et al. 2009; Ling et al. 2010; Pillepich et al. 2014). The kinematics of Nyx-2 suggest that it is related to Nyx and may be tidal debris from a separate passage of the same satellite. We approach this conclusion cautiously only because the network scores for Nyx-2 stars are generally lower than those for Nyx, and its detection is thus not as robust.

To study the chemical abundances of each component, we use the cross-match of stars provided by Sanders & Das (2018), covering the spectroscopic catalogs APOGEE (Majewski et al. 2017), LAMOST (Deng et al. 2012), RAVE (Steinmetz et al. 2006), GALAH (de Silva et al. 2015), Gaia-ESO (Gilmore et al. 2012), and SEGUE (Yanny et al. 2009). Out of the 9379 (22,296) stars in the high-purity (canonical) sample, we find cross-matches for 4255 (10,309). We histogram each star’s metallicity Z in Figure 7. A star is associated with a given component if the mixture analysis finds that it has a probability greater than 95% (75%) of belonging to the high-purity (canonical) sample.¹⁴ The Halo has a mean at $Z \sim -1.1$ in the high-purity sample, with a large dispersion and a long tail toward more metal-poor stars. Fewer stars are associated with the Halo in the canonical sample as compared to the high-purity sample (less than 3% versus 10%), but the metallicities are largely consistent between the two. Enceladus also has a large dispersion in Z and peaks at $Z \sim -0.9$ in the high-purity sample and $Z \sim -1.0$ in the canonical sample. Nyx stars, however, have a much narrower Z distribution, peaking at $Z \sim -0.6$ ($Z \sim -0.4$) in the high-purity (canonical) sample. In

¹⁴ We lower the cut for the canonical sample because fewer stars are exclusively associated to a single Gaussian with a probability higher than 95% as a result of the larger number of components.

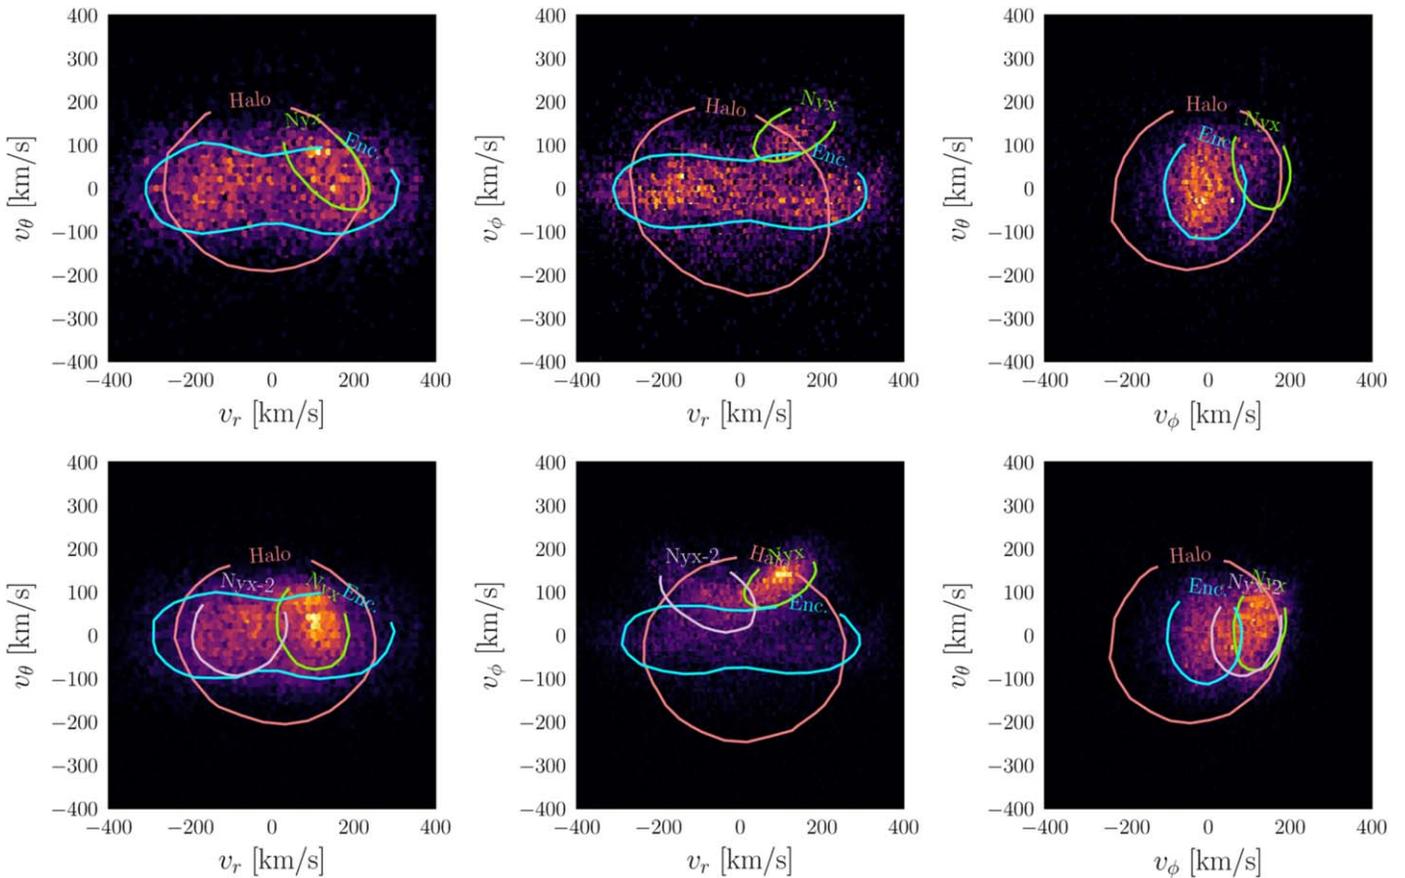


Figure 5. Kinematic distributions of stars in the Gaia DR2 catalog that have measured radial velocities and fall within Galactocentric radii $r \in [6.5, 9.5]$ kpc and vertical distances $|z| < 3$ kpc of the midplane. These stars have been identified as accreted by the neural network developed in Ostdiek et al. (2020); the top row shows the distributions for the high-purity sample ($S > 0.95$) while the bottom row shows the distributions for the canonical sample ($S > 0.85$). The two-dimensional distributions are shown for the Galactocentric velocity coordinates v_r , v_θ , and v_ϕ . In the high-purity sample, a Gaussian mixture study recovers a Halo component with large dispersion (pink), Gaia Enceladus (blue), and Nyx (green). These same components are also identified in the canonical sample, which additionally includes a separate prograde stream, which we refer to as Nyx-2 (purple). Note that “Halo” refers to the remaining accreted stars in the stellar halo that are not individually resolved by the mixture analysis.

both the canonical and high-purity samples, the distributions of the different components are consistent. Interestingly, the distribution of Nyx-2 in the canonical sample matches that of Nyx, strengthening the argument that they are debris from the same progenitor. Spectroscopic follow-ups will be crucial to more deeply understand the origin of Nyx and Nyx-2.

In Figure 8, we provide a scatter plot of the maximum vertical distance from the midplane, z_{\max} , versus the eccentricity, e , of the different components as obtained from a simple calculation of their orbits. These parameters are evaluated when running the stellar orbits back in time 1 Gyr over 1000 steps, using the package `gala` (Price-Whelan 2017). To take the measurement errors into account, we resample each star 100 times from the position and velocity Gaussians, and average their eccentricities and z_{\max} . Figure 8 only shows 100 random stars from each component to not overwhelm the plot. Full distributions of the eccentricities, z_{\max} , apocenters, and pericenters of each component are provided in the Appendix, Figures A11–A13.

Enceladus stars have largely eccentric orbits ($e > 0.8$), confirming that the satellite progenitor was on a highly radial orbit (Deason et al. 2018). Similarly, Halo stars have large eccentricities and extend up to ~ 10 kpc above the plane, on average higher than Enceladus. The 1D histograms of the eccentricities show that Enceladus stars are peaked toward

larger eccentricities (Figures A12 and A13). Nyx stars extend to $z_{\max} \sim 0.2\text{--}4$ kpc above the plane at eccentricities $e \sim 0.4\text{--}0.8$ —the latter are higher than what is expected of thick-disk stars (Li & Zhao 2017).

4. More Subtle Kinematic Structures

In Section 3, we discussed the three most dominant structures found in the high-purity and canonical catalogs. To see if there is any evidence for non-Gaussian structures, we now use the algorithm DBSCAN¹⁵ (Pedregosa et al. 2011). This clustering algorithm is designed to identify connected overdensities; we will apply it to the 3D velocities in Galactocentric cylindrical coordinates, not spherical as in Section 3, which will allow for easier comparison with previous studies. Note that DBSCAN does not require any assumptions regarding the shape of the underlying distributions and therefore can identify non-Gaussian structures. This is particularly important when hunting for small stream-like features. DBSCAN depends on the minimum number of stars in a group, n_{\min} , and the compactness parameter, ϵ .

It is important to take into account the errors on the positions and velocities of the stars while using DBSCAN. To do so, we

¹⁵ <https://scikit-learn.org/stable/modules/clustering.html>

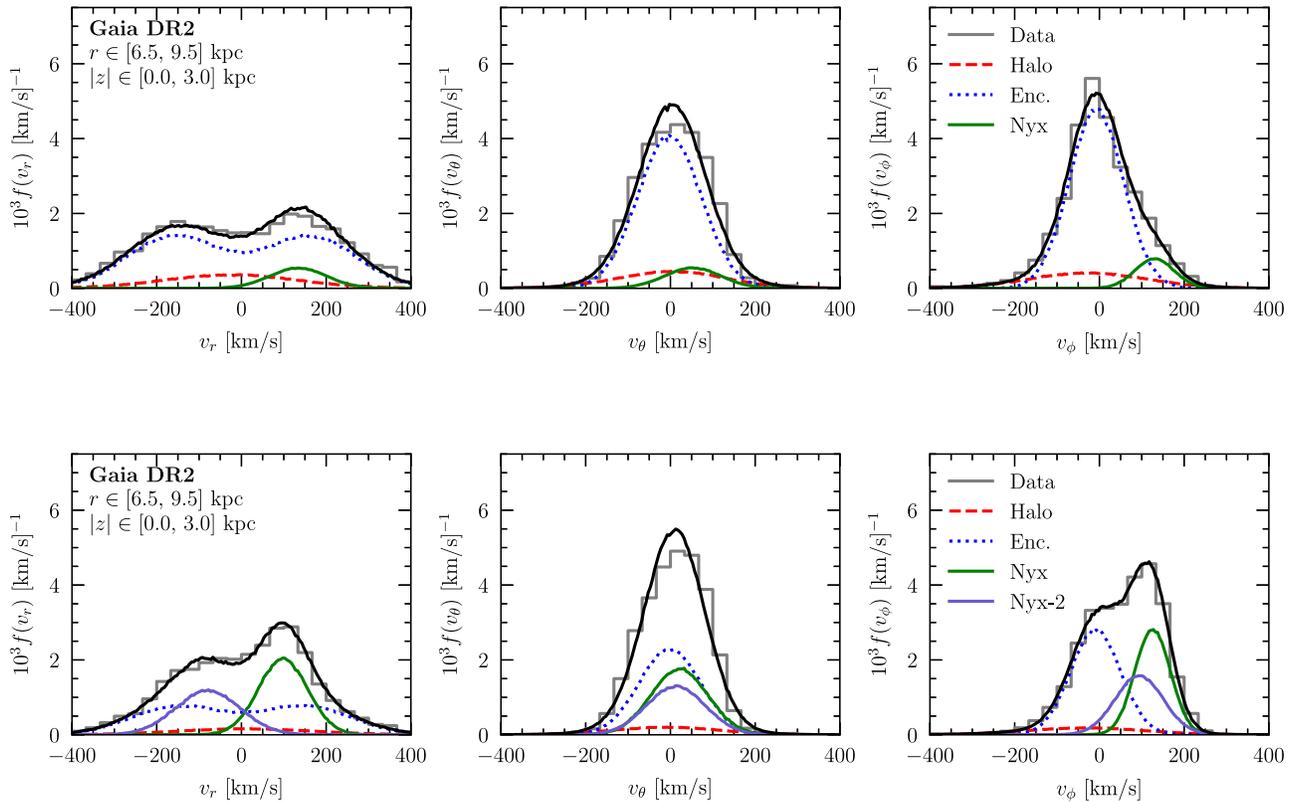


Figure 6. Best-fit velocity distributions for all stars with network scores $S > 0.95$ (top row) and $S > 0.85$ (bottom row), which have measured radial velocities and fall within Galactocentric radii $r \in [6.5, 9.5]$ kpc and vertical distances $|z| < 3$ kpc of the midplane. The distributions are shown for the Halo (red dashed), Enceladus (blue dotted), and Nyx (green solid) components. The total distribution of the model is shown in solid black and can be compared to the data (gray histogram). In the canonical sample ($S > 0.85$), we also identify a second prograde stream, called Nyx-2. A plot of the model residuals for the high-purity sample is provided in Figure A9.

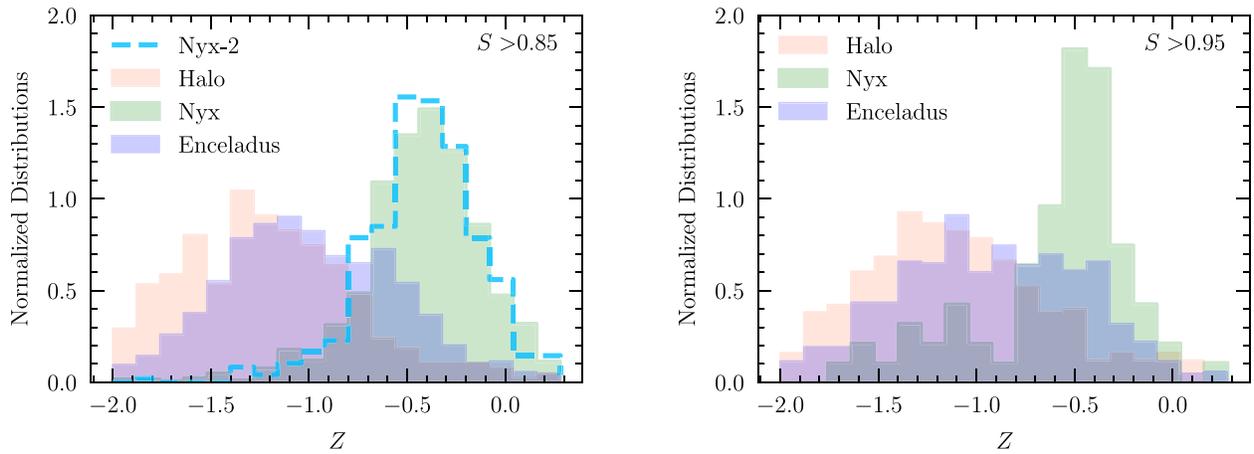


Figure 7. The metallicity Z distributions for the Halo, Nyx, and Enceladus components in the canonical (left) and high-purity (right) samples. The use of Z here should not be confused with z , the vertical distance above the plane of the Milky Way. We also show the distribution for Nyx-2, which is only identified in the canonical sample. A star is associated with any one of these populations if it has a $>95\%$ ($>85\%$) probability of belonging to the respective Gaussian component for the high-purity (canonical) sample. All distributions are normalized to unity. Note that the relative normalizations between components reflect the fractions within our sample. Metallicities are taken from Sanders & Das (2018).

resample the stars over their errors and run DBSCAN 100 times, identifying the groups for each independent run. In Figure 9, we show the resulting groups for $n_{\min} = 5$ and $\epsilon = 0.4$ for one such realization, which returned nine groups. Note that one of these groups (shown in pale orange) is attributed to the combination of Gaia Enceladus and Nyx, as their distributions overlap, and they are thus not separated by this algorithm.

Each realization of DBSCAN is prone to statistical fluctuations as we vary the stars within their measurement errors. To find the most robust structures, we save the centers of each of the groups identified in each of the 100 resamplings and then run DBSCAN over these 100 realizations to find clusters of group centers that are present in more than 70 realizations, with $\epsilon = 0.3$ and $n_{\min} = 50$. We find four stream candidates that

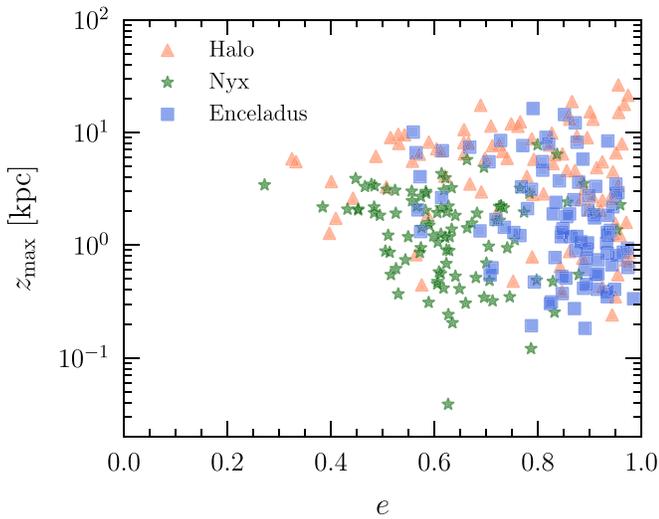


Figure 8. Scatter plot of the maximum vertical distance versus eccentricity for Nyx, Halo, and Gaia Enceladus stars in the high-purity sample. We only show a random selection of 100 stars for each component. Nyx stars are typically concentrated in the disk plane. However, their eccentricities are larger than expected for thick-disk stars (Li & Zhao 2017).

pass these cuts for the high-purity sample; see Table 2. Interestingly, when running the same algorithm on the canonical sample, the same three groups pass these requirements.

The stream that is the most robust—in that it is found by all 100 realizations—is an overdensity at $(v_R, v_\phi, v_z) = (21, 141, -286) \text{ km s}^{-1}$, shown in the realization of Figure 9 as the stars in yellow (labeled as Group I). This overdensity can be matched to the Helmi stream (Helmi et al. 1999),¹⁶ which has been reported in Gaia DR2 by several studies, including Gaia Collaboration et al. (2018b; Koppelman et al. 2018). We identify ~ 20 stars as belonging to this overdensity, all with $v_z < 0 \text{ km s}^{-1}$. We test the IDs of the core Helmi stream stars provided in Koppelman et al. (2019) against our framework to find their associated network scores. Of the 40 IDs given, 37 pass our parallax error cut. Of these, 10 have $S < 0.5$ and 27 have high scores of $S > 0.5$, with 22 having scores $S > 0.9$, meaning that the neural network is confident they are accreted. We do not reconstruct the second (smaller) known cluster of the Helmi stream at $v_z > 0 \text{ km s}^{-1}$.

The other three robust streams are located at $(v_R, v_\phi, v_z) = (-154, -377, -63) \text{ km s}^{-1}$, $(v_R, v_\phi, v_z) = (-195, -237, 164) \text{ km s}^{-1}$, and $(212, -229, 161) \text{ km s}^{-1}$ with small dispersions,¹⁷ as listed in Table 2. These structures fall near several velocity clusters identified in Koppelman et al. (2018) and may be related.

5. Conclusions

Stars in the Milky Way galaxy can be divided into two components: those that were born within the Galaxy and those that were accreted. The phase-space distribution of accreted stars provides a crucial handle for understanding how the Galaxy evolved by revealing the imprints of satellite mergers.

¹⁶ It also overlaps the S2 stream in Myeong et al. (2018a), which is believed to be related to the Helmi stream.

¹⁷ For Groups II and III, we only provide the maximum dispersions. In some cases, the group has only two stars, making it difficult to define a meaningful velocity dispersion.

This approach requires distinguishing the population of accreted stars from their in situ counterparts, a task that becomes increasingly challenging near the Galactic midplane where disk stars comprise $\sim 99\%$ of all stars. This motivated the work of Ostdiek et al. (2020), where we developed a deep neural-network-based approach that allows us to build a catalog of accreted stars from Gaia DR2 data. Although the network provides a catalog of all well-measured Gaia stars, regardless of whether or not they have a line-of-sight velocity measurement, this paper provides the first analysis of the 4.8 million star subset that includes the full 6D information and that fall within $r \in [6.5, 9.5] \text{ kpc}$ and $|z| < 3 \text{ kpc}$.

Our primary goal is to identify and analyze structures in 3D velocity space. As a first step, we perform a Gaussian mixture analysis to break down the high-purity sample ($S > 0.95$) into its most significant contributions: Gaia Enceladus, Nyx, and the Halo. We find that Enceladus is highly radial and comprises the vast majority of accreted stars in this region of the sky. These results are consistent with previous studies of Enceladus, which characterized its properties farther from the disk plane (Belokurov et al. 2018; Deason et al. 2018; Helmi et al. 2018; Lancaster et al. 2019; Necib et al. 2019a). Nyx is a new stream identified by this analysis. It is prograde and comprises nearly 9% of the high-purity sample, making it one of the most significant streams to be discovered to date near the Sun. Properties and a discussion of the potential origin of Nyx as a merging dwarf galaxy are explored in greater detail in Necib et al. (2020). The “Halo” is essentially the remaining group of accreted stars that cannot be further subdivided by the mixture analysis.

We also repeated the analysis on all accreted stars with network scores $S > 0.85$. This canonical sample is considerably larger in size than the high-purity one, but likely has more contamination from disk stars. We again recover Enceladus and Nyx in this sample. Additionally, we find evidence for another prograde stream, Nyx-2, with roughly the same rotational speed as Nyx and equal—but opposite—radial velocity. Nyx-2 comprises $\sim 22\%$ of the canonical sample, comparable to the Nyx fraction in this sample. Similarities in the kinematics and metallicities between Nyx and Nyx-2 suggest that they may be related to the same progenitor.

We also attempt to reconstruct non-Gaussian velocity structures using the DBSCAN algorithm. We locate four additional streams using this method, one of which coincides with the well-studied Helmi stream (Helmi et al. 1999). The Helmi stream is by far the most robust of the four, consisting of ~ 20 stars and identified over all repeated iterations that account for uncertainties in the stellar velocity measurements. The other three candidate streams consist of fewer stars and are recovered $\sim 70\%$ of the time over repeated iterations of DBSCAN. These streams are all retrograde and may be associated with overdensities identified in Koppelman et al. (2018).

The analysis presented here demonstrates the power derived from combining advancements in data quality, numerical simulations, and data analysis techniques. The fact that the catalog reproduces known structures such as Gaia Enceladus and the Helmi stream validates the utility of this approach. However, the science case extends much further than simple validation. Indeed, the new catalog greatly improves our understanding of the stellar distribution in the ROI studied. In particular, it clearly demonstrates that Enceladus extends down into the Galactic plane. It also unearths evidence for a

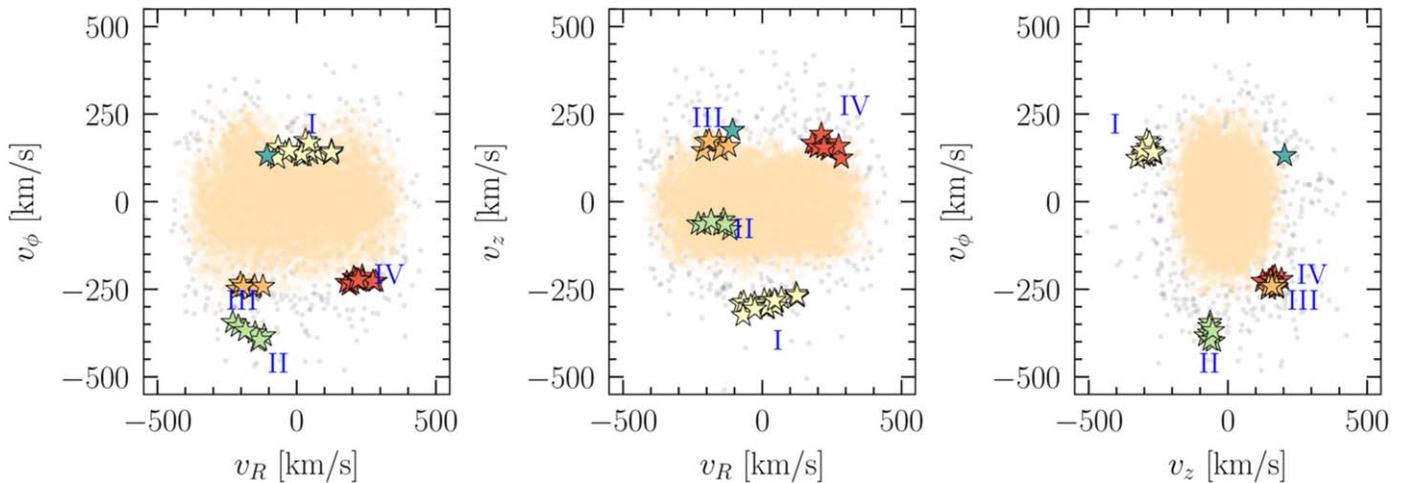


Figure 9. Results of running the DBSCAN clustering algorithm on the high-purity ($S > 0.95$) accreted stars in the catalog for a single realization. The clusters are shown in Galactocentric cylindrical coordinates (v_R , v_ϕ , v_z). The colored stars indicate groups of stars that are associated kinematically; the gray dots are outliers identified by the method. The large pale orange cluster of points corresponds to the Enceladus and Nyx structures, which DBSCAN associates together. We repeat this procedure 100 times, sampling over the errors in stellar position and velocity. Groups I–IV are the most robust as they are identified in over 70 of the 100 iterations. The green star without an associated number is an example of a cluster that is not present through more than 70 iterations.

Table 2
Centers of the Velocity Clusters Identified by DBSCAN after 100 Iterations Scanning over Position and Velocity Errors of the Stars

ID	v_R (km s $^{-1}$)	v_ϕ (km s $^{-1}$)	v_z (km s $^{-1}$)	σ_R (km s $^{-1}$)	σ_ϕ (km s $^{-1}$)	σ_z (km s $^{-1}$)	N_{stars}	Frequency
I	21	141	−286	40–77	8–26	8–24	10–20	100%
II	−154	−377	−63	63	34	18	3–19	95%
III	−195	−237	164	41	24	24	3–17	78%
IV	212	−229	161	9–35	7–25	1–37	7–17	74%

Note. For each cluster, we provide the following: ID number, mean velocities in Galactocentric cylindrical coordinates, range of dispersions of stars that belong in the group through error resampling, number of stars, and its frequency across 100 reruns. We only list groups that have occurred more than 70% of the time. The velocities of Group I overlap with those of the Helmi stream (Helmi et al. 1999), while Groups II, III, and IV may be related to streams identified in Koppelman et al. (2018). The same groups appear when analyzing either the high-purity or the canonical sample.

significant new stellar stream (Nyx), a potentially related counterpart (Nyx-2), and three other smaller candidate streams.

We have improved over previous approaches as a result of two main factors. The first is due to the statistical benefit of an increased overall size of the accreted stellar sample. Additionally, having used a deep network that is only trained on phase space allows us to derive a high-purity sample of accreted stars without imposing strong cuts on circular velocity or metallicity, as is typically done. This reduces the intrinsic bias that results from such cuts. We note that this paper has only scratched the surface, and in particular it will be very interesting to investigate what structures can be identified in the rest of our Gaia DR2 catalog that does not include line-of-sight velocities.

Understanding the mergers that contributed stellar debris in our neighborhood of the Milky Way has the potential to provide an empirical determination of the local dark matter distribution, which is also built up from mergers (Herzog-Arbeitman et al. 2018; Necib et al. 2019b), and as such, is expected to include remnant structures. The recent discovery of Enceladus, for example, motivates extending the Standard Halo Model of dark matter to (at least) a two-component model that includes both an isotropic halo and debris flow (Necib et al. 2019a). By clearly demonstrating that Enceladus extends into the disk plane, the results of this work confirm that dark matter debris from this merger likely contributes in the solar

neighborhood. The discovery of Nyx near the solar position may suggest the presence of a corresponding dark matter stream or disk. Coupling this catalog with cosmological simulations will be essential in refining our understanding of the local dark matter phase-space distribution, and its implications for direct detection experiments.

We thank G. Brova, P. Hopkins, E. Kirby, R. Sanderson, and A. Wetzel for helpful discussions. This work was performed in part at Aspen Center for Physics, which is supported by National Science Foundation grant PHY-1607611. 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.

L.N. is supported by the DOE under Award Number DESC0011632, the Sherman Fairchild fellowship, the California Presidential fellowship, and a Carnegie Fellowship in Theoretical Astrophysics. B.O. and T.C. are supported by the US Department of Energy under grant No. DE-SC0011640. M.L. is supported by the DOE under award number DESC0007968 and the Cottrell Scholar Program through the Research Corporation for Science Advancement. M.F. is supported by the Zuckerman STEM Leadership Program and

in part by the DOE under grant No. DE-SC0011640. S.G.K. is 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.

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.

Software: This analysis made use of *Astropy* (Price-Whelan et al. 2018), *Galpy* (Bovy 2015), *Matplotlib* (Hunter 2007), *NumPy* (van der Walt et al. 2011), and *Scikit-Learn* (Pedregosa et al. 2011). The neural network used for tagging the accreted stars was implemented in *Keras* (Chollet 2015) with the *TensorFlow* backend (Abadi et al. 2015). The network was trained using Adam (Kingma & Ba 2014) to minimize the binary crossentropy loss.

Note Added

As this work was being completed, the paper by Borsato et al. (2020) became available. Using DBSCAN applied to the integrals of motion, Borsato et al. (2020) found the groups of stars we also identify in Table 2, along with four more that do not pass our selection criteria.

Appendix

In this Appendix, we show two-dimensional histograms of stars in v_y - $\sqrt{v_x^2 + v_z^2}$ space with $S < 0.05$ and $S \in [0.3, 0.5]$ in Figure A1. We find that the scores roughly track the different components of the sample, with $S < 0.05$ largely identifying the thin disk, and $S \in [0.3, 0.5]$ largely identifying the thick disk. Thick-disk stars are harder to discern from the thin disk and the (largely accreted) halo and thus receive midrange scores. Remember that the neutral network is only identifying in situ and accreted stars—the fact that the network scores roughly correlate with physical components of the Galaxy is an added bonus. In Figures A2 and A3, we provide the corner plots of the fractions of each component in the Gaussian mixture model analysis for the high-purity and canonical samples, respectively. In Figures A4 and A5, we show the best-fit parameters of Nyx in the high-purity (canonical) sample. The best-fit parameters of Nyx-2 in the canonical sample are shown in Figure A6. Similar treatment of Enceladus is shown in Figures A7 and A8 for the high-purity and canonical samples, respectively. The residual of the Gaussian mixture analysis of the high-purity sample is shown in Figure A9. In Figure A10, we provide the spatial distribution of Enceladus stars and focus particularly on how they extend down to the Galactic midplane. In Figures A11–A15, we provide the orbital properties of Nyx, Enceladus, and Halo for the high-purity sample, as well as stars with scores $S < 0.05$ and $S \in [0.3, 0.5]$ for reference.

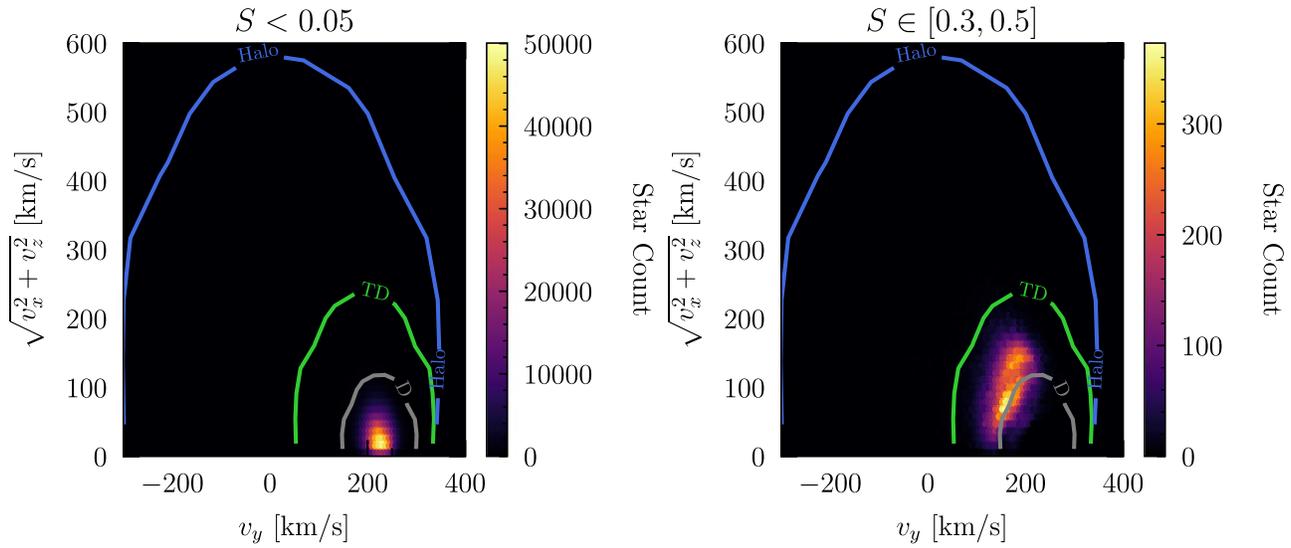


Figure A1. Toomre plots for stars with scores $S < 0.05$ (left) and $S \in [0.3, 0.5]$ (right) along with the 3σ velocity contours for the thin disk (D), thick disk (TD), and stellar halo (Halo), following Bensby et al. (2003).

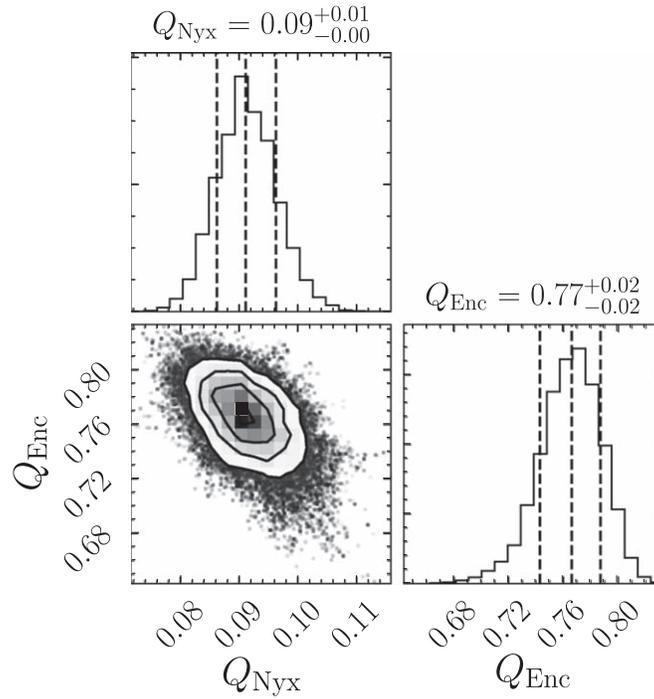


Figure A2. Corner plot of the `emcee` run for the fractions of the different components in the high-purity sample.

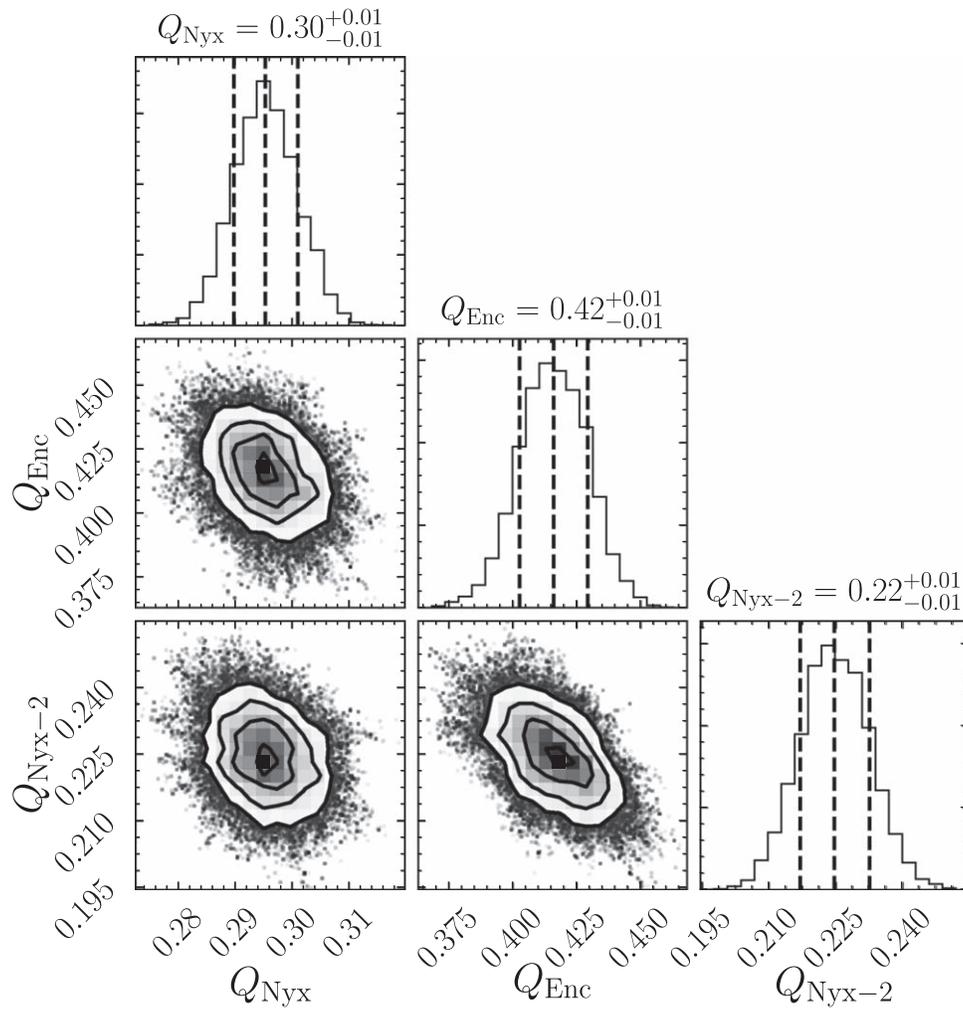


Figure A3. Corner plot of the emcee run for the fractions of the different components in the canonical sample.

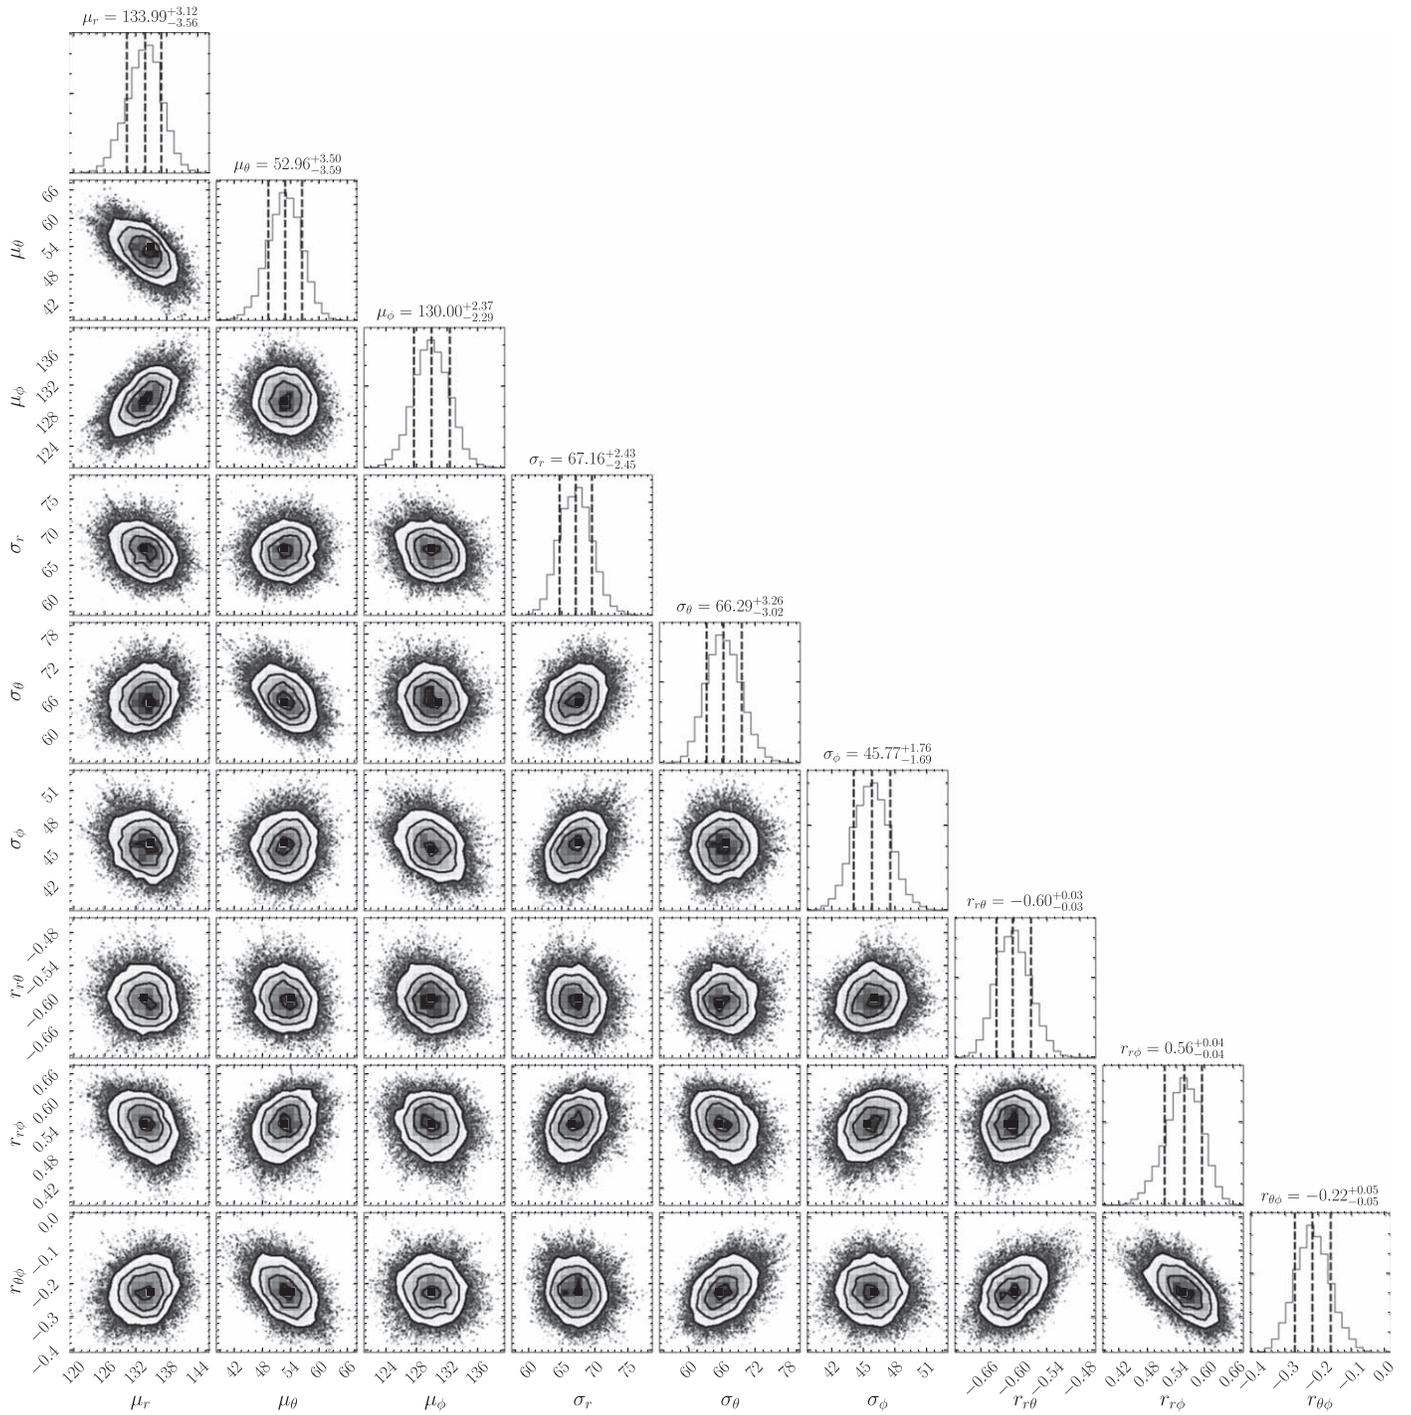


Figure A4. Corner plot of the emcee run for Nyx in the high-purity sample.

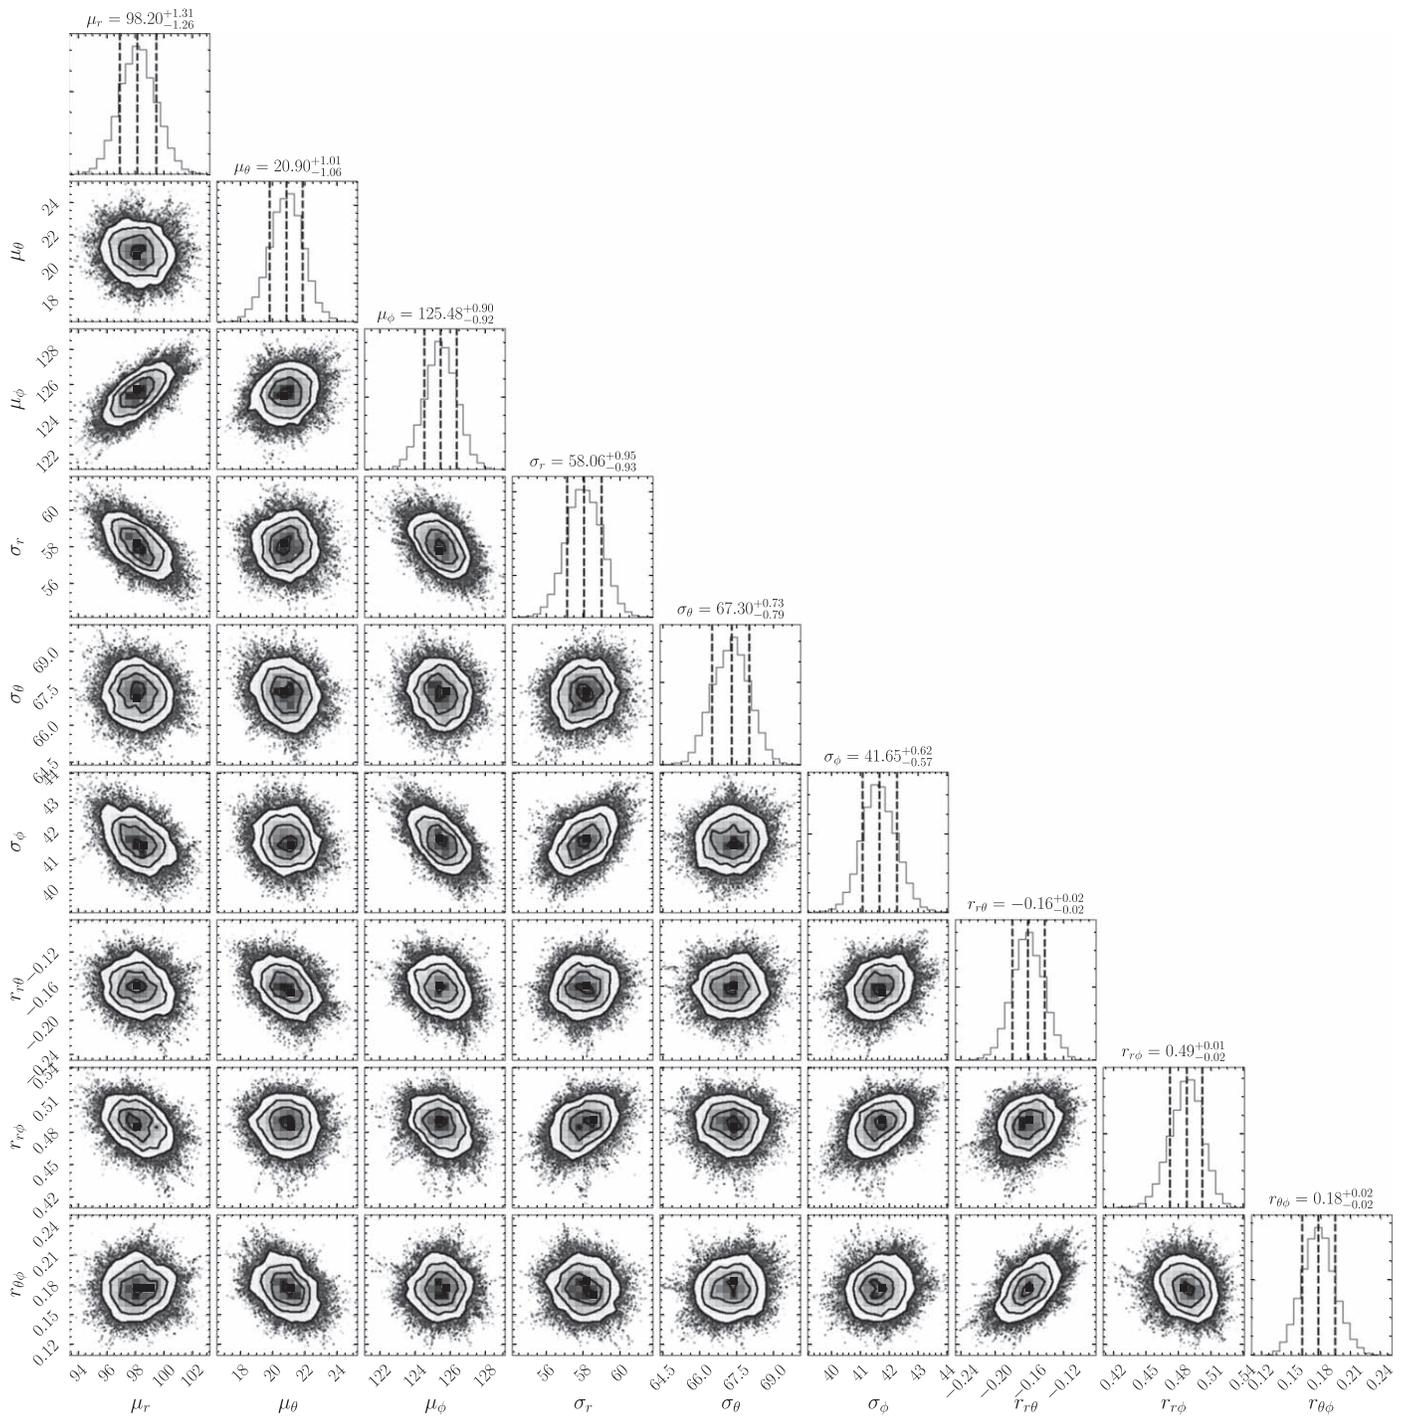


Figure A5. Corner plot of the emcee run for Nyx in the canonical sample.

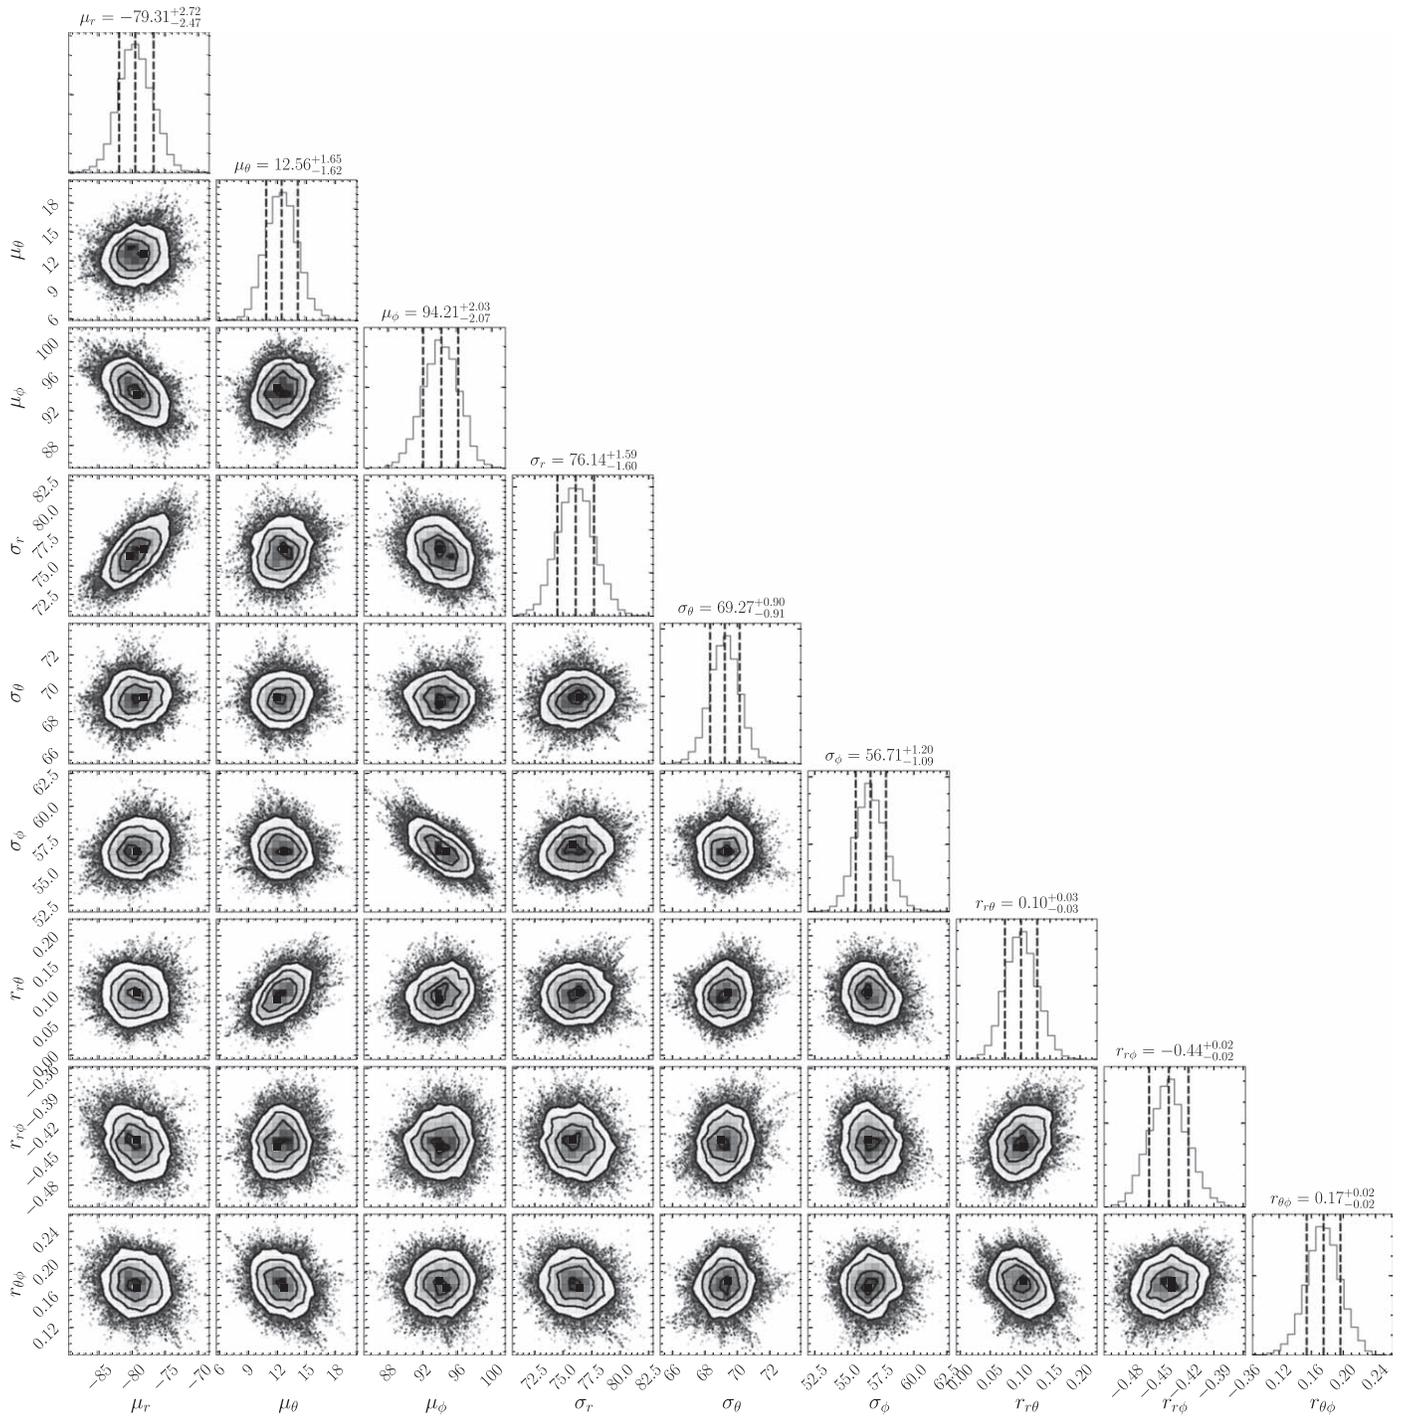


Figure A6. Corner plot of the emcee run for Nyx-2 in the canonical sample.

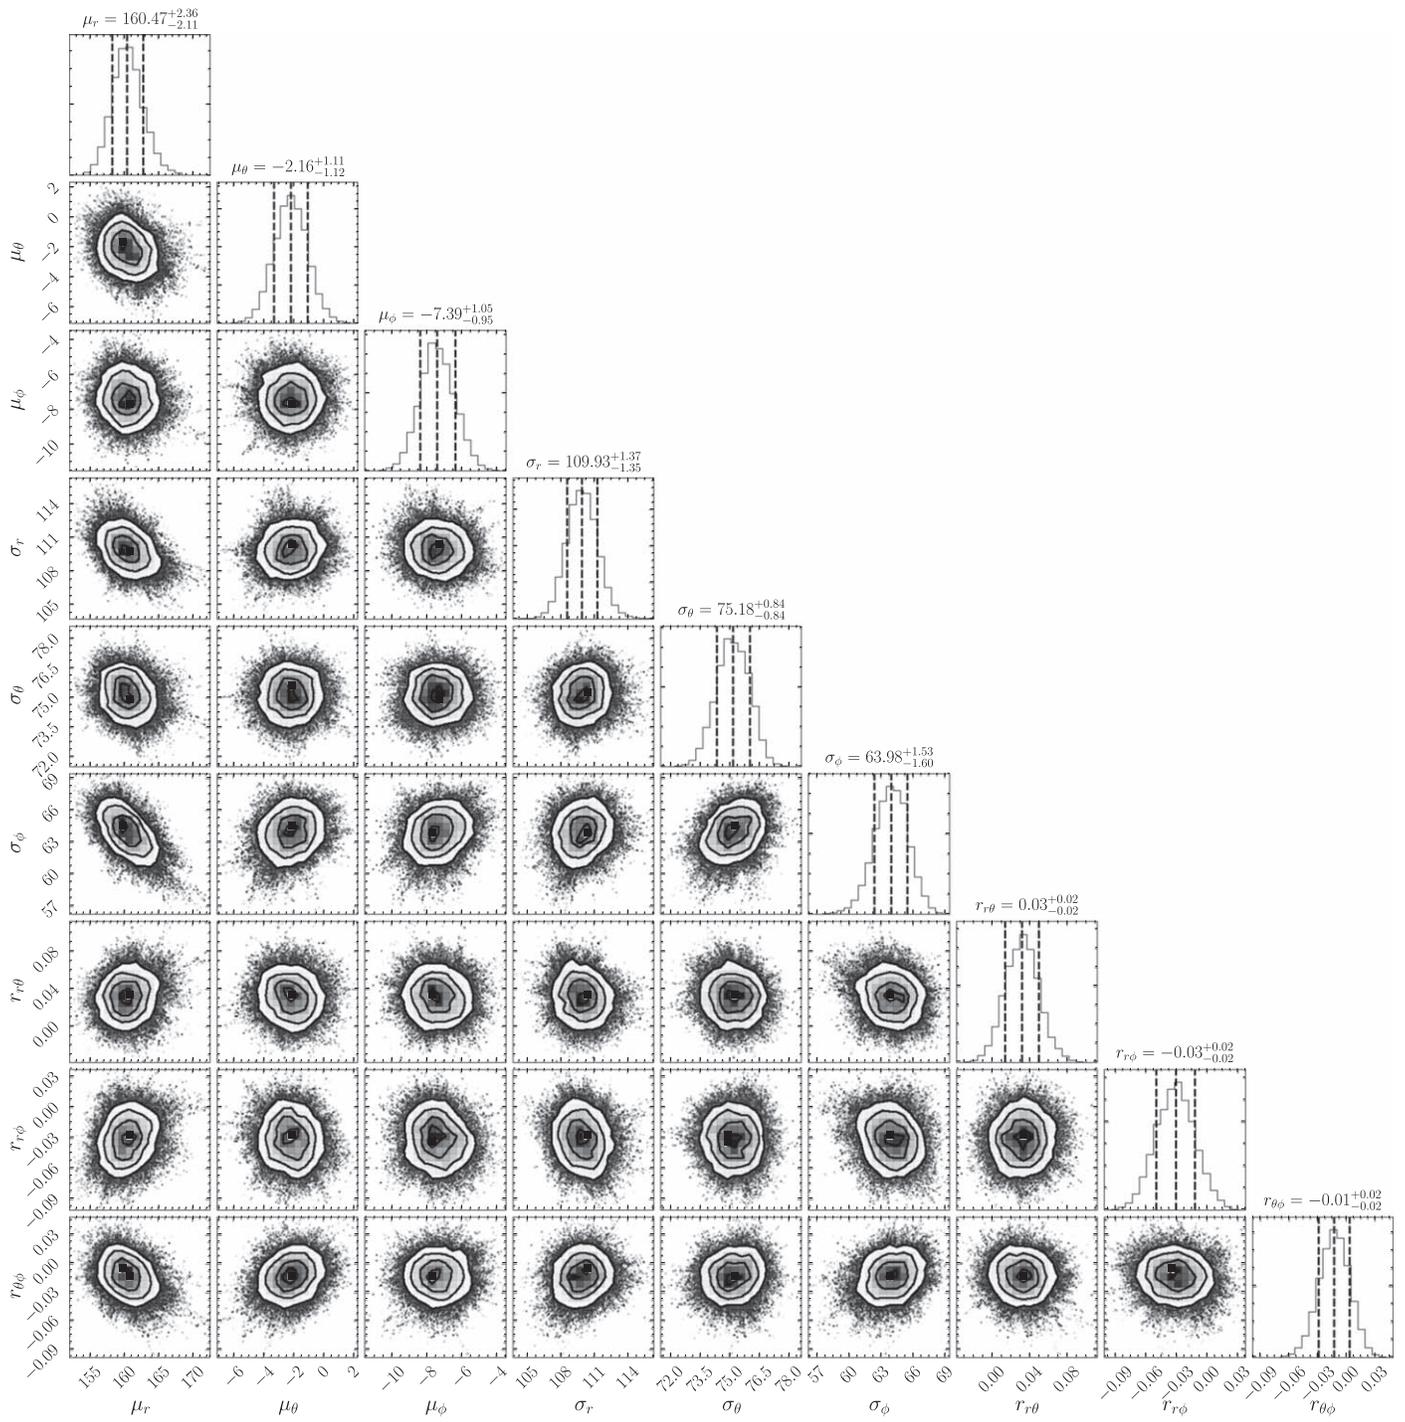


Figure A7. Corner plot of the emcee run for Enceladus in the high-purity sample.

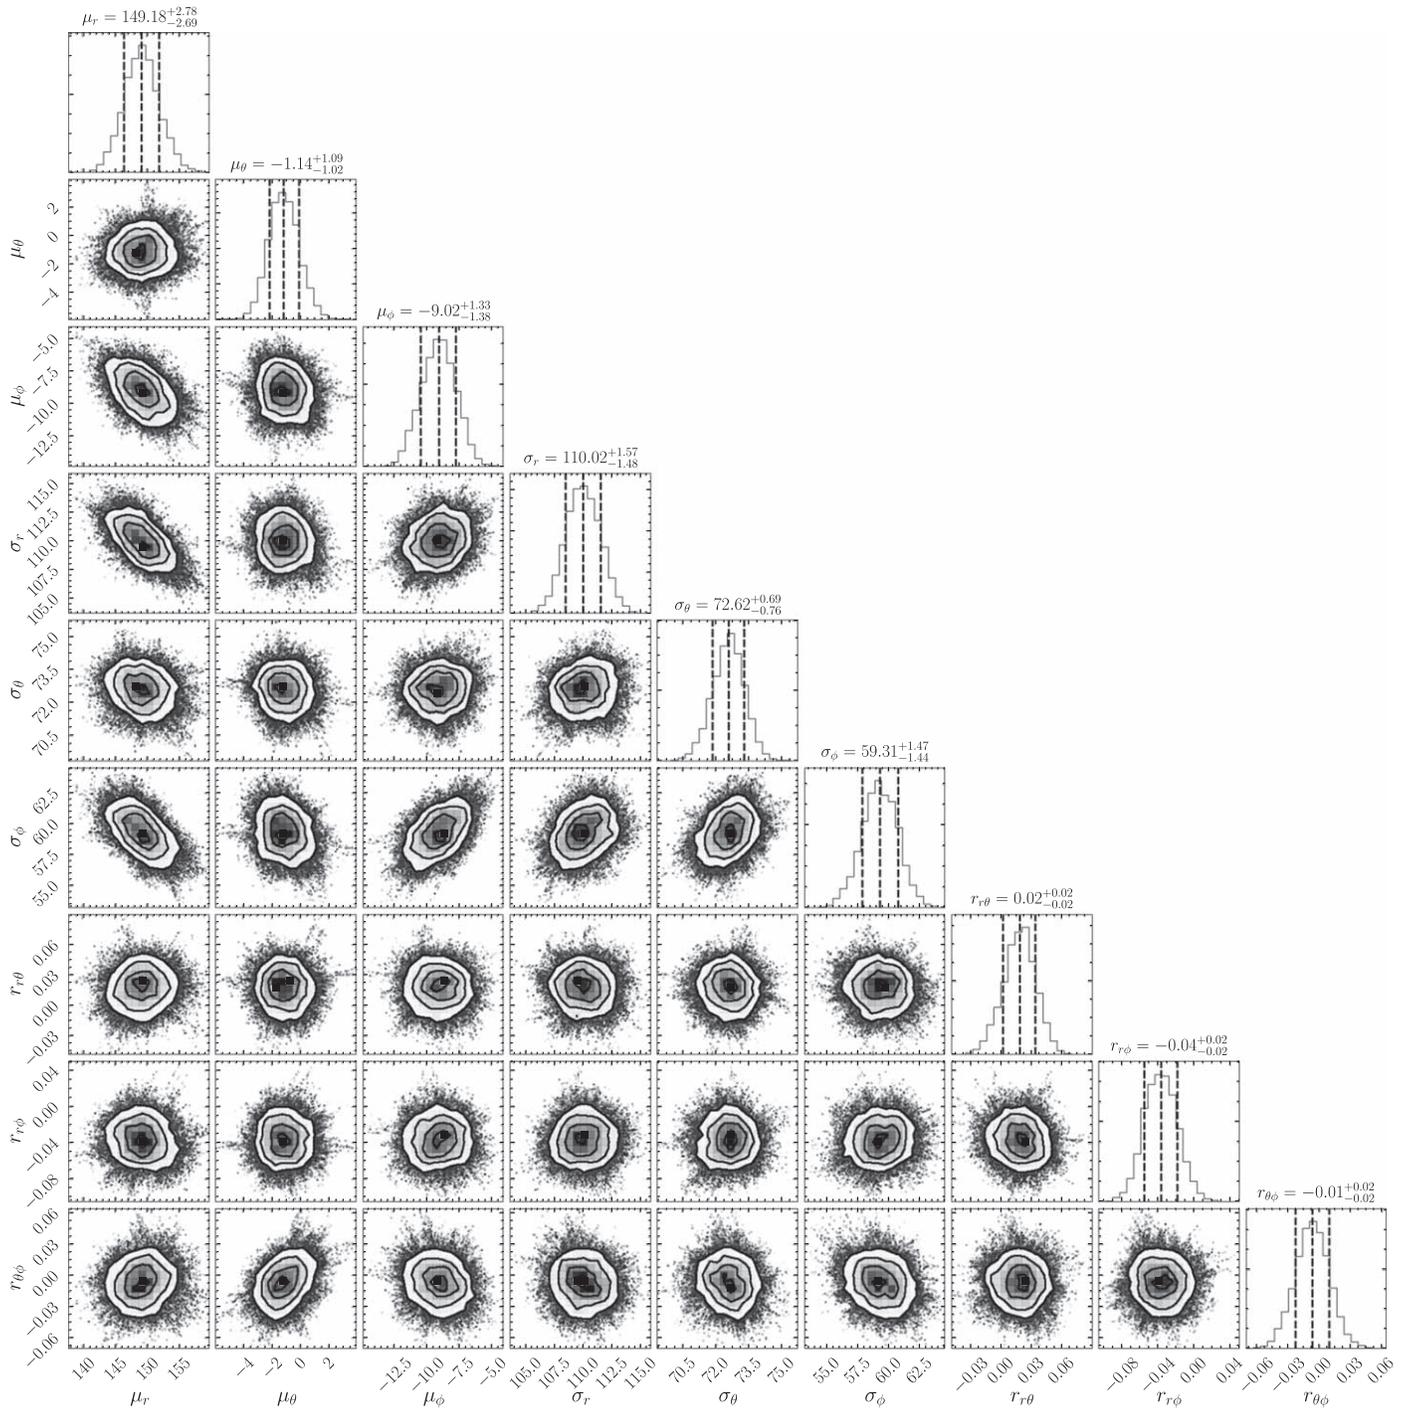


Figure A8. Corner plot of the emcee run for Enceladus in the canonical sample.

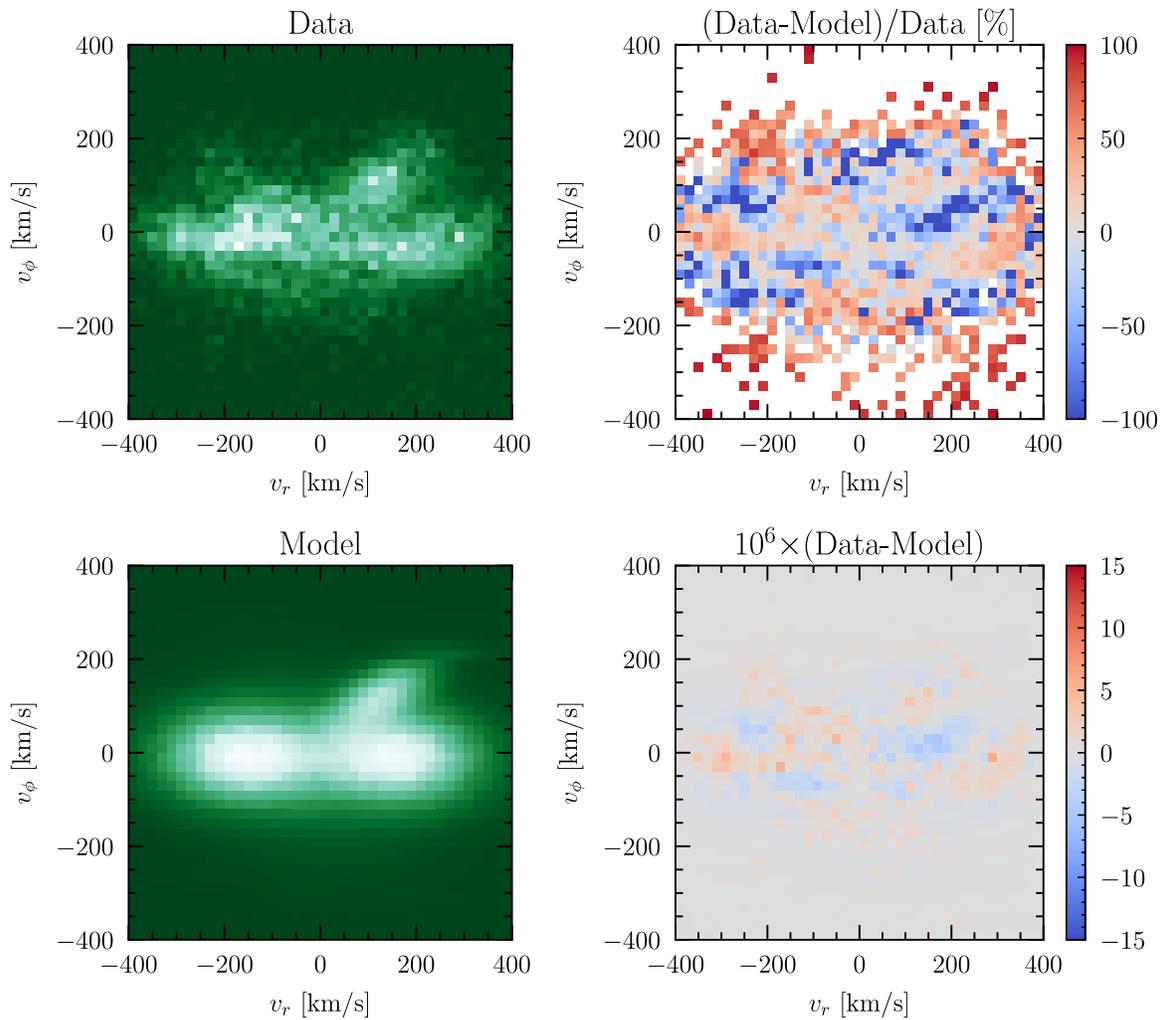


Figure A9. Residuals for the high-purity sample. The top-left panel shows the data in the $v_r - v_\phi$ plane. The bottom left panel shows the model prediction. The right panels show the residuals (Data-Model); these are normalized to the data count in the top panel.

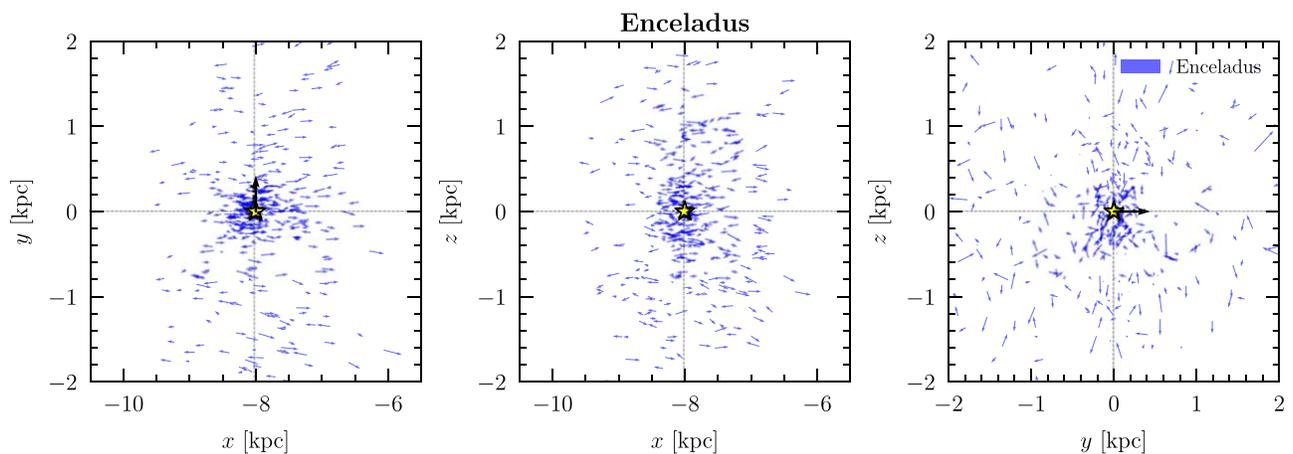


Figure A10. Spatial distribution of the velocity vectors for the stars associated with Gaia Enceladus in the high-purity sample. The Sun is located at $(x, y, z) = (-8, 0, 0)$ kpc, and the black arrow indicates its velocity. The sample of Enceladus stars has been subsampled by a factor of 5 for ease of viewing.

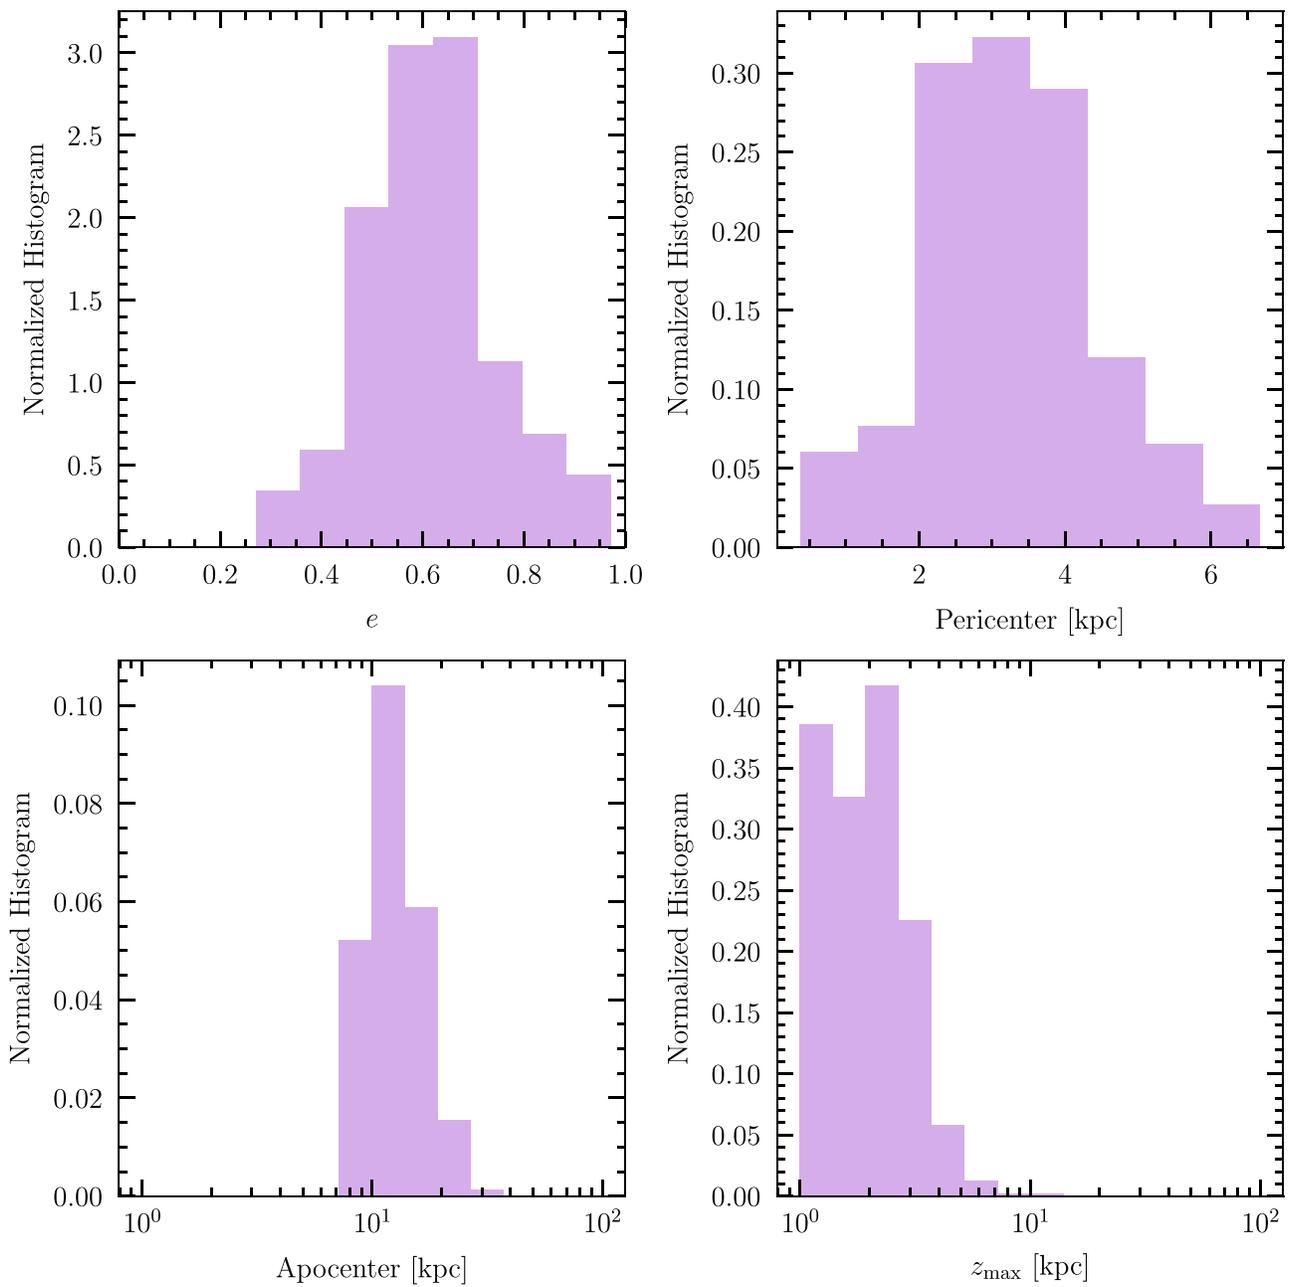


Figure A11. Orbital properties of stars associated with Nyx, calculated as described in the main text. We show distributions for the eccentricities (top left), pericenters (top right), apocenters (bottom left), and z_{\max} (bottom right). The orbits were evolved back 1 Gyr over 1000 steps using `gala` (Price-Whelan 2017), assuming the default Milky Way potential in Bovy (2015).

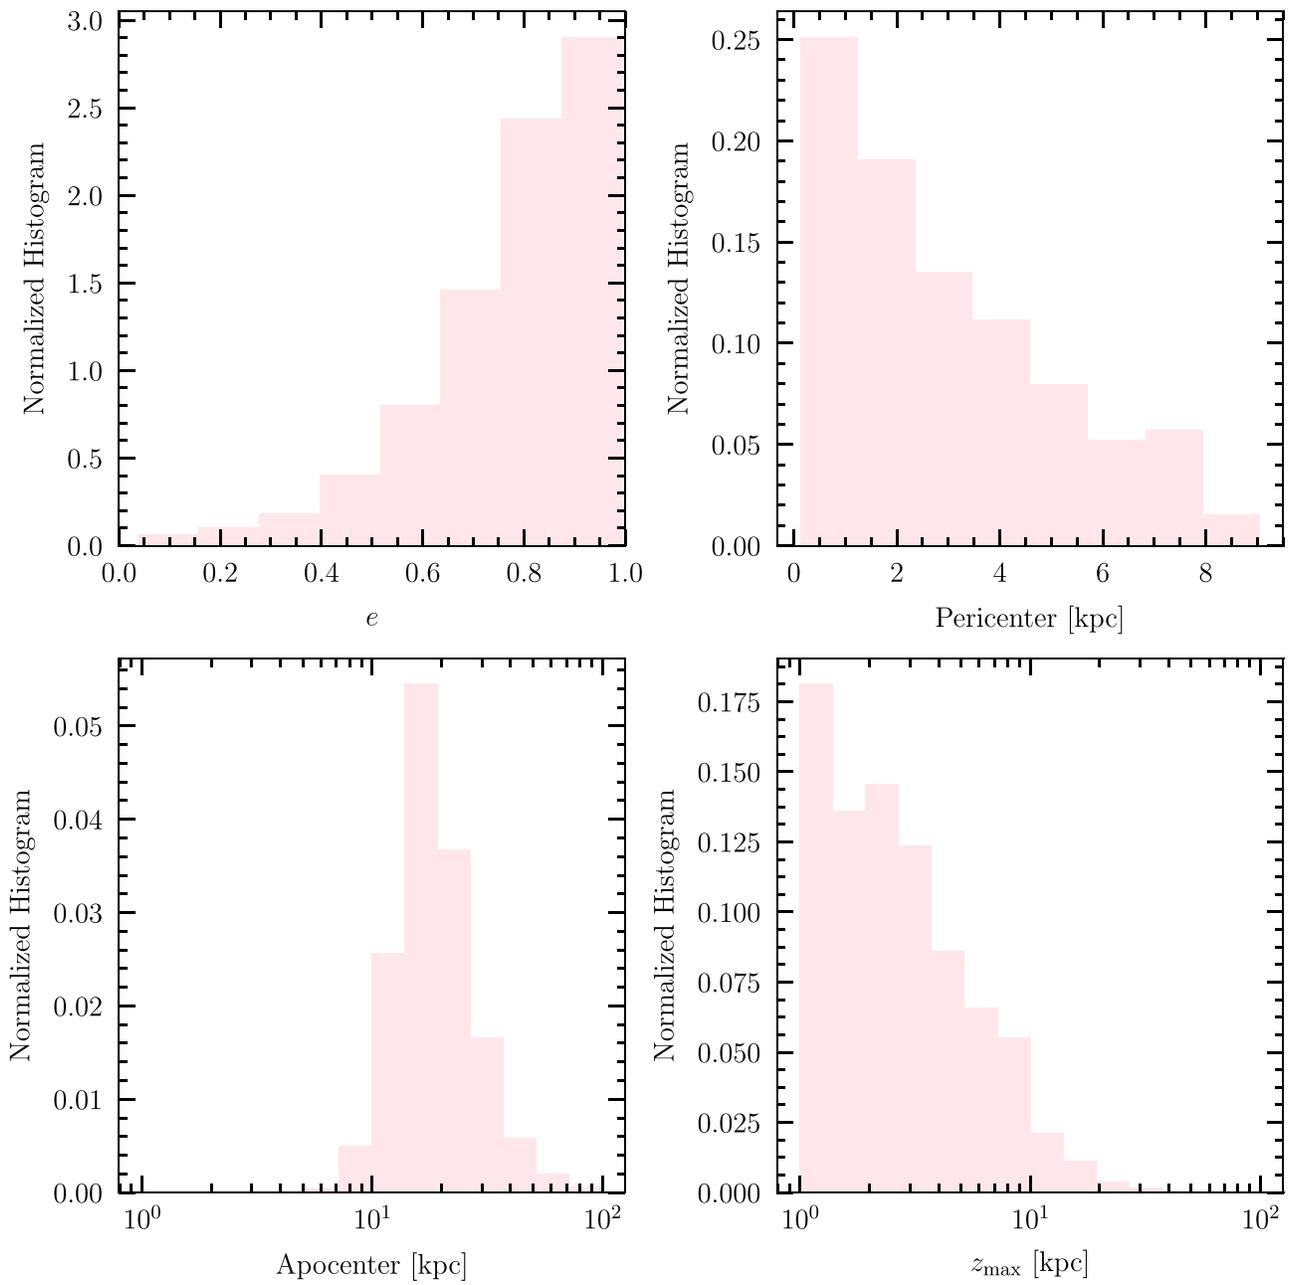


Figure A12. Same as Figure A11, except for the halo stars.

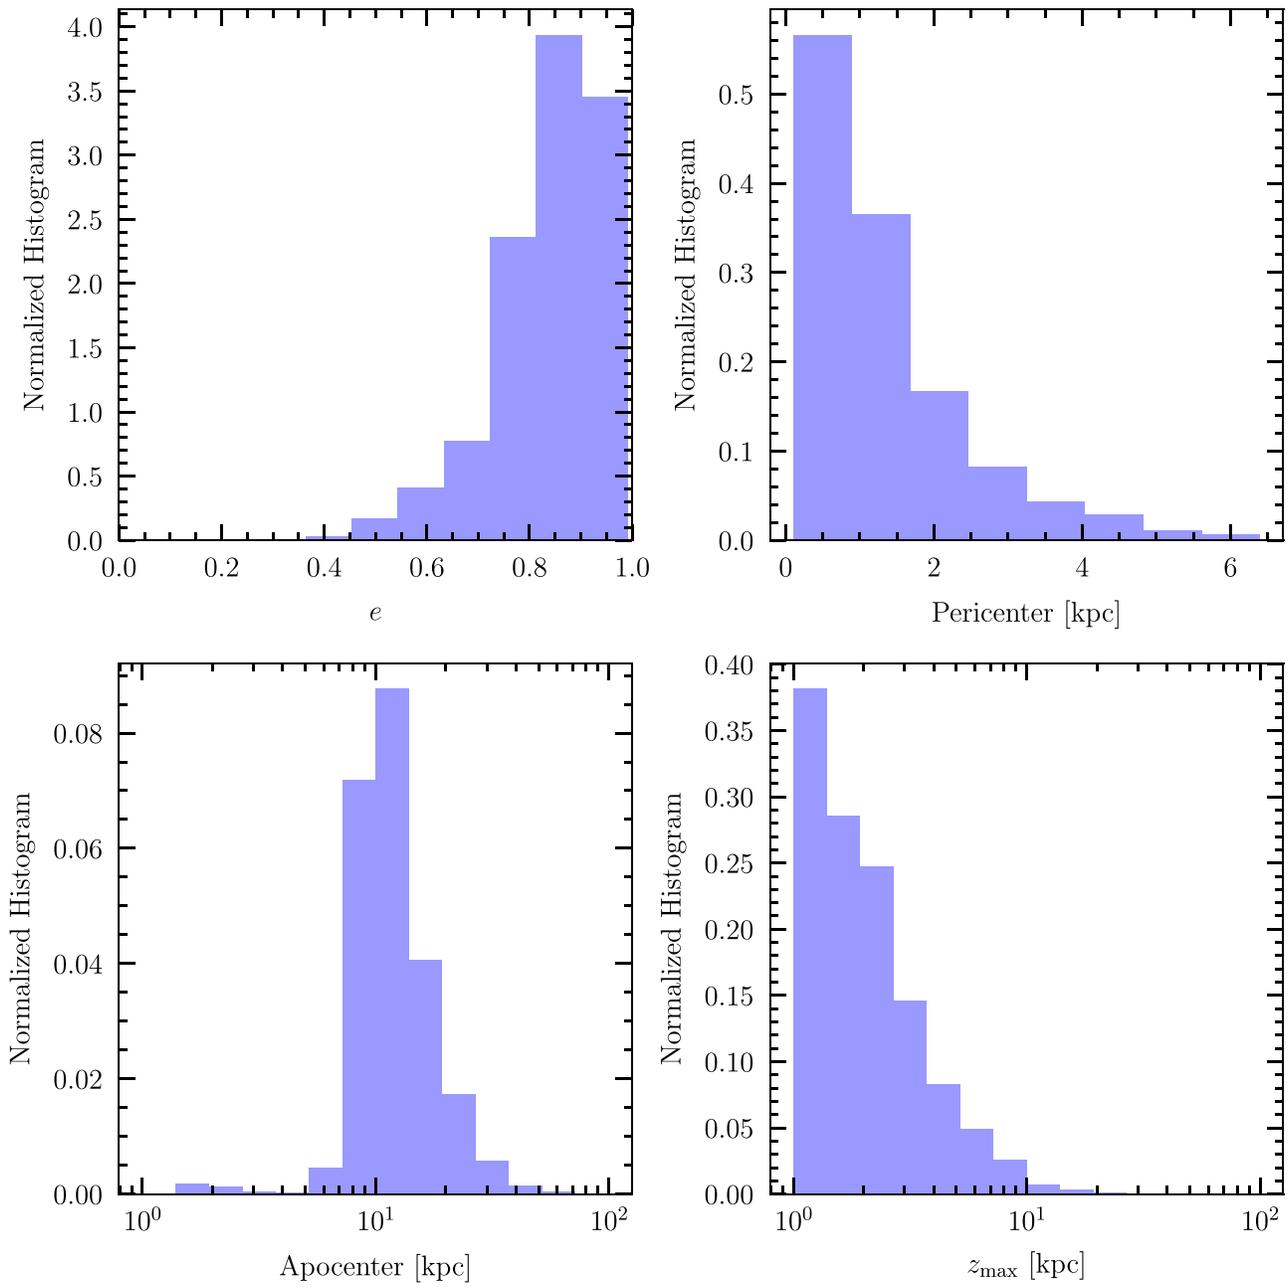


Figure A13. Same as Figure A11, except for the Enceladus stars.

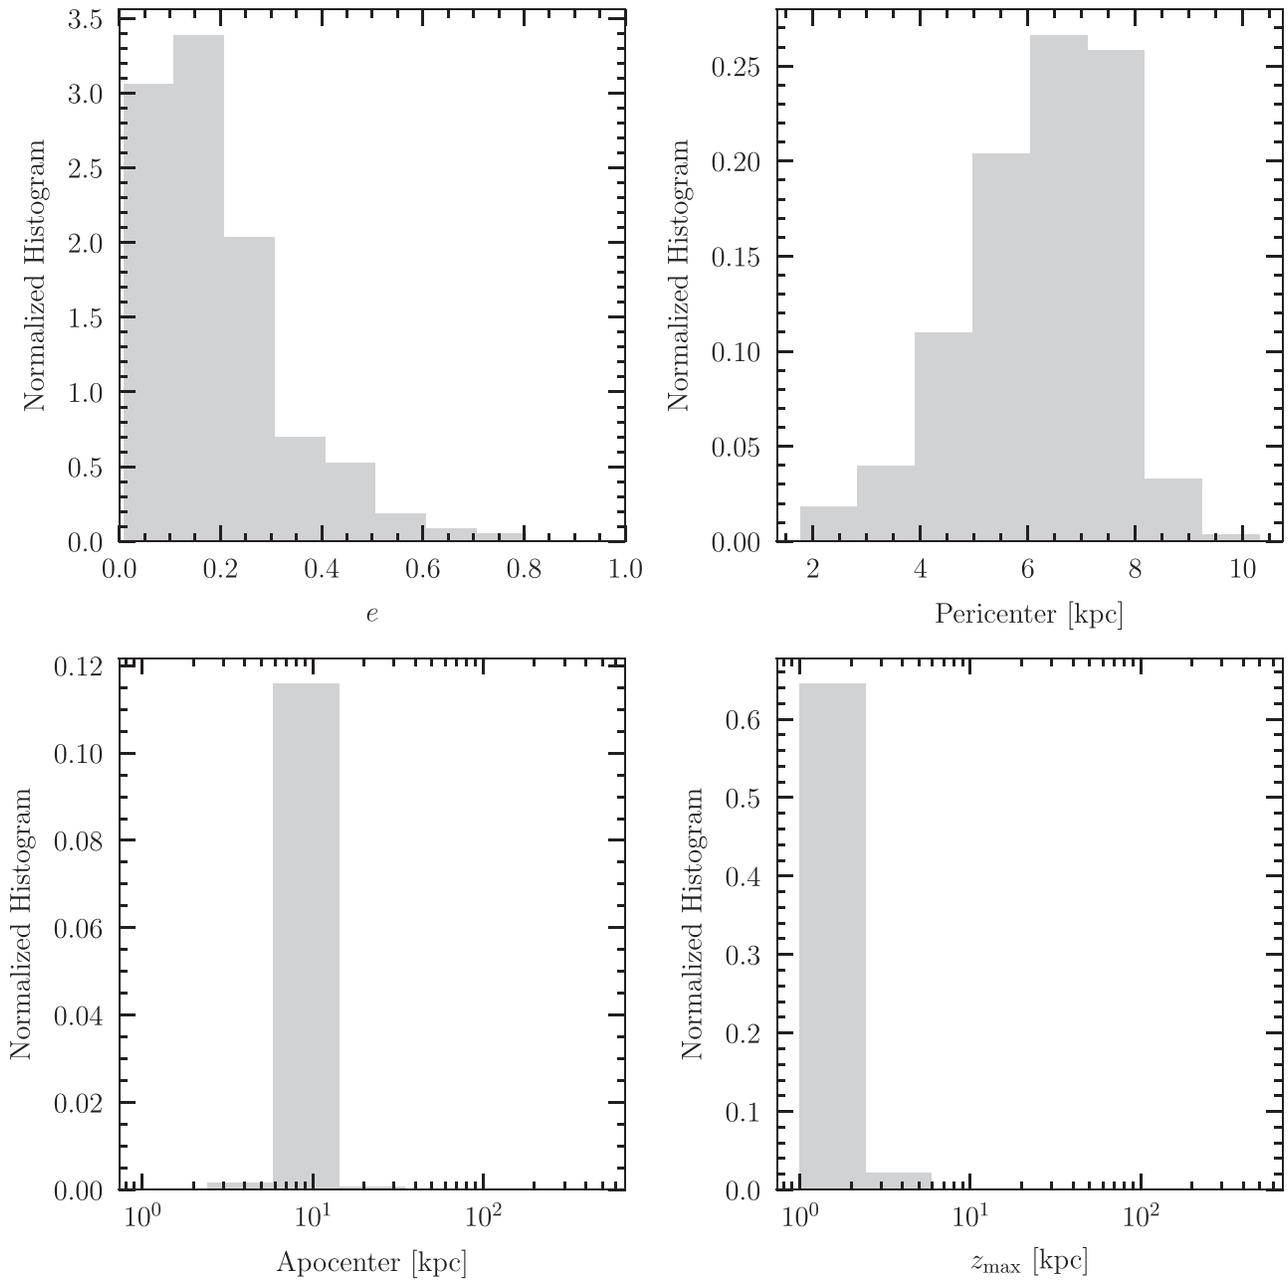


Figure A14. Same as Figure A11, except for stars with network scores $S < 0.05$.

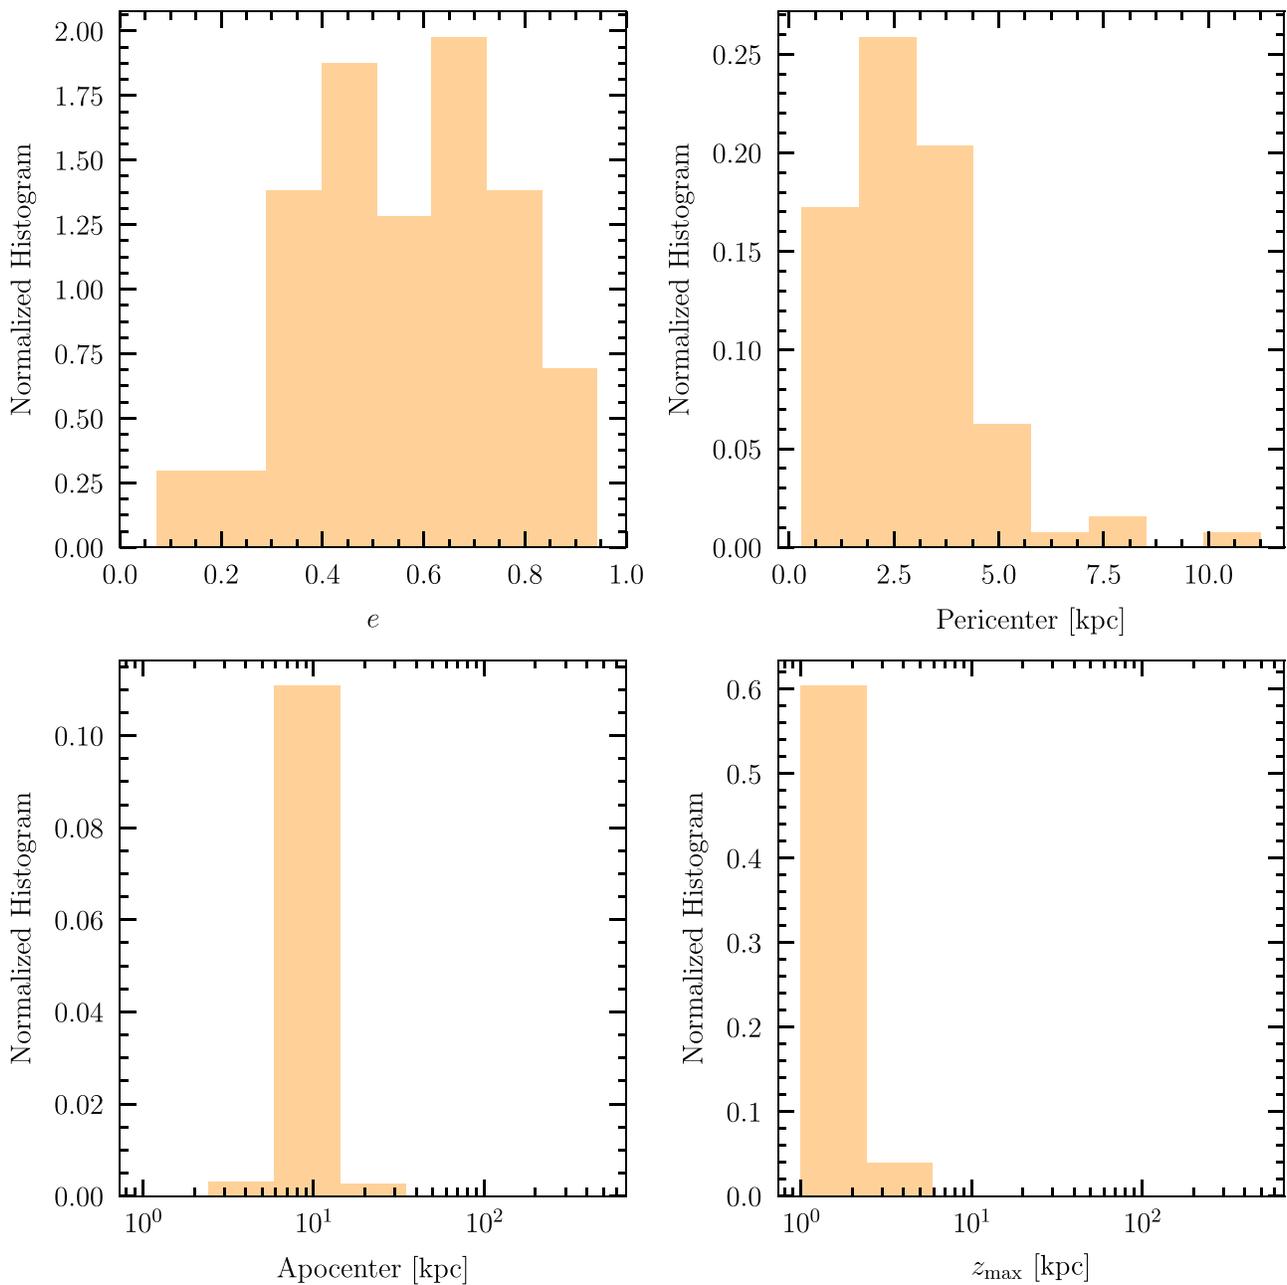


Figure A15. Same as Figure A11, except for stars with network scores $S \in [0.3, 0.5]$.

ORCID iDs

Lina Necib  <https://orcid.org/0000-0003-2806-1414>
 Bryan Ostdiek  <https://orcid.org/0000-0002-0376-6461>
 Shea Garrison-Kimmel  <https://orcid.org/0000-0002-4655-8128>

References

- Abadi, M., Agarwal, A., Barham, P., et al. 2015, TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems, www.tensorflow.org
- Abadi, M. G., Navarro, J. F., Steinmetz, M., & Eke, V. R. 2003, *ApJ*, **597**, 21
- An, D., & Beers, T. C. 2020, *ApJ*, **897**, 39
- An, D., Beers, T. C., Johnson, J. A., et al. 2013, *ApJ*, **763**, 65
- An, D., Beers, T. C., Santucci, R. M., et al. 2015, *ApJL*, **813**, L28
- Beers, T. C., Carollo, D., Ivezić, Ž, et al. 2012, *ApJ*, **746**, 34
- Beers, T. C., Placco, V. M., Carollo, D., et al. 2017, *ApJ*, **835**, 81
- Belokurov, V., Erkal, D., Evans, N. W., Koposov, S. E., & Deason, A. J. 2018, *MNRAS*, **478**, 611
- Bensby, T., Feltzing, S., & Lundström, I. 2003, *A&A*, **410**, 527
- Bernard, E. J., Ferguson, A. M. N., Schlafly, E. F., et al. 2014, *MNRAS: Lett.*, **443**, L84
- Borsato, N. W., Martell, S. L., & Simpson, J. D. 2020, *MNRAS*, **492**, 1370
- Bovy, J. 2015, *ApJS*, **216**, 29
- Brown, A. G. A., Vallenari, A., Prusti, T., et al. 2018, *A&A*, **616**, A1
- Bullock, J. S., & Johnston, K. V. 2005, *ApJ*, **635**, 931
- Carollo, D., Beers, T. C., Chiba, M., et al. 2010, *ApJ*, **712**, 692
- Carollo, D., Beers, T. C., Lee, Y. S., et al. 2007, *Natur*, **450**, 1020
- Chiba, M., & Beers, T. C. 2000, *AJ*, **119**, 2843
- Chollet, F. 2015, Keras: Deep learning library for theano and tensorflow, 7, T1, <https://keras.io>
- de Jong, J. T. A., Yanny, B., Rix, H.-W., et al. 2010, *ApJ*, **714**, 663
- Deason, A. J., Belokurov, V., Koposov, S. E., & Lancaster, L. 2018, *ApJL*, **862**, L1
- Deng, L.-C., Newberg, H. J., Liu, C., et al. 2012, *RAA*, **12**, 735

- de Silva, G. M., Freeman, K. C., Bland-Hawthorn, J., et al. 2015, *MNRAS*, **449**, 2604
- Font, A. S., Johnston, K. V., Bullock, J. S., & Robertson, B. E. 2006, *ApJ*, **646**, 886
- Font, A. S., McCarthy, I. G., Crain, R. A., et al. 2011, *MNRAS*, **416**, 2802
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, **125**, 306
- Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2018a, *A&A*, **616**, A1
- Gaia Collaboration, Helmi, A., van Leeuwen, F., et al. 2018b, *A&A*, **616**, A12
- Gilmore, G., Randich, S., Asplund, M., et al. 2012, *Msngr*, **147**, 25
- Grillmair, C. J., & Dionatos, O. 2006, *ApJL*, **643**, L17
- Helmi, A., Babusiaux, C., Koppelman, H. H., et al. 2018, *Natur*, **563**, 85
- Helmi, A., & de Zeeuw, P. T. 2000, *MNRAS*, **319**, 657
- Helmi, A., Navarro, J. F., Nordstrom, B., et al. 2006, *MNRAS*, **365**, 1309
- Helmi, A., Veljanoski, J., Breddels, M. A., Tian, H., & Sales, L. V. 2017, *A&A*, **598**, A58
- Helmi, A., & White, S. D. M. 1999, *MNRAS*, **307**, 495
- Helmi, A., White, S. D. M., de Zeeuw, P. T., & Zhao, H. 1999, *Natur*, **402**, 53
- Herzog-Arbeitman, J., Lisanti, M., Madau, P., & Necib, L. 2018, *PhRvL*, **120**, 041102
- Hopkins, P. F. 2015, *MNRAS*, **450**, 53
- Hopkins, P. F., Wetzel, A., Kereš, D., et al. 2018, *MNRAS*, **480**, 800
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Ibata, R., Lewis, G. F., Irwin, M., Totten, E., & Quinn, T. R. 2001, *ApJ*, **551**, 294
- Ibata, R. A., Bellazzini, M., Malhan, K., Martin, N., & Bianchini, P. 2019a, *NatAs*, **3**, 667
- Ibata, R. A., Gilmore, G., & Irwin, M. J. 1994, *Natur*, **370**, 194
- Ibata, R. A., Malhan, K., & Martin, N. F. 2019b, *ApJ*, **872**, 152
- Ivezic, Z., Goldston, J., Finlator, K., et al. 2000, *AJ*, **120**, 963
- Johnston, K. V. 1998, *ApJ*, **495**, 297
- Johnston, K. V., Hernquist, L., & Bolte, M. 1996, *ApJ*, **465**, 278
- Johnston, K. V., Spergel, D. N., & Hernquist, L. 1995, *ApJ*, **451**, 598
- Jurić, M., Ivezić, Ž., Brooks, A., et al. 2008, *ApJ*, **673**, 864
- Katz, D., Sartoretti, P., Cropper, M., et al. 2019, *A&A*, **622**, A205
- Kepley, A. A., Morrison, H. L., Helmi, A., et al. 2007, *AJ*, **134**, 1579
- Kingma, D. P., & Ba, J. 2014, arXiv:1412.6980
- Klement, R., Rix, H. W., Flynn, C., et al. 2009, *ApJ*, **698**, 865
- Koppelman, H., Helmi, A., & Veljanoski, J. 2018, *ApJL*, **860**, L11
- Koppelman, H. H., Helmi, A., Massari, D., Roelenga, S., & Bastian, U. 2019, *A&A*, **625**, A5
- Kunder, A., Kordopatis, G., Steinmetz, M., et al. 2017, *AJ*, **153**, 75
- Lake, G. 1989, *AJ*, **98**, 1554
- Lancaster, L., Koposov, S. E., Belokurov, V., Evans, N. W., & Deason, A. J. 2019, *MNRAS*, **486**, 378
- Li, C., & Zhao, G. 2017, *ApJ*, **850**, 25
- Ling, F. S., Nezri, E., Athanassoula, E., & Teyssier, R. 2010, *JCAP*, **2010**, 012
- Lisanti, M., & Spergel, D. N. 2012, *PDU*, **1**, 155
- Lisanti, M., Spergel, D. N., & Madau, P. 2015, *ApJ*, **807**, 14
- Majewski, S. R., Schiavon, R. P., Frinchaboy, P. M., et al. 2017, *AJ*, **154**, 94
- Malhan, K., & Ibata, R. A. 2018, *MNRAS*, **477**, 4063
- Meingast, S., Alves, J., & Fűrncranz, V. 2019, *A&A*, **622**, L13
- Myeong, G. C., Evans, N. W., Belokurov, V., Amorisco, N. C., & Koposov, S. E. 2018a, *MNRAS*, **475**, 1537
- Myeong, G. C., Evans, N. W., Belokurov, V., Sanders, J. L., & Koposov, S. E. 2018b, *ApJL*, **863**, L28
- Myeong, G. C., Vasiliev, E., Iorio, G., Evans, N. W., & Belokurov, V. 2019, *MNRAS*, **488**, 1235
- Necib, L., Bryan, O., Lisanti, M., et al. 2020, *NatAs*
- Necib, L., Lisanti, M., & Belokurov, V. 2019a, *ApJ*, **874**, 3
- Necib, L., Lisanti, M., Garrison-Kimmel, S., et al. 2019b, *ApJ*, **883**, 27
- Newberg, H. J., & Carlin, J. L. (ed.) 2016, *Tidal Streams in the Local Group and Beyond*, Vol. 420 (1st edn.; Berlin: Springer)
- Nissen, P. E., & Schuster, W. J. 2010, *A&A*, **511**, L10
- Odenkirchen, M., Grebel, E. K., Rockosi, C. M., et al. 2001, *ApJL*, **548**, L165
- Ostdiek, B., Necib, L., Cohen, T., et al. 2019, *Cataloging Accreted Stars within Gaia DR2 Using Deep Learning*, Zenodo, doi:10.5281/zenodo.3579379
- Ostdiek, B., Necib, L., Cohen, T., et al. 2020, *A&A*, **636**, A75
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *J. Mach. Learn. Res.*, **12**, 2825
- Pillepich, A., Kuhlen, M., Guedes, J., & Madau, P. 2014, *ApJ*, **784**, 161
- Posti, L., Helmi, A., Veljanoski, J., & Breddels, M. A. 2018, *A&A*, **615**, A70
- Price-Whelan, A. M. 2017, *JOSS*, **2**, 388
- Price-Whelan, A. M., Sipőcz, B. M., Günther, H. M., et al. 2018, *AJ*, **156**, 123
- Purcell, C. W., Bullock, J. S., & Kaplinghat, M. 2009, *ApJ*, **703**, 2275
- Quinn, P. J., & Goodman, J. 1986, *ApJ*, **309**, 472
- Re Fiorentin, P., Helmi, A., Lattanzi, M. G., & Spagna, A. 2005, *A&A*, **439**, 551
- Read, J. I., Lake, G., Agertz, O., & Debattista, V. P. 2008, *MNRAS*, **389**, 1041
- Read, J. I., Mayer, L., Brooks, A. M., Governato, F., & Lake, G. 2009, *MNRAS*, **397**, 44
- Rix, H.-W., & Bovy, J. 2013, *A&ARv*, **21**, 61
- Robertson, B., Bullock, J. S., Font, A. S., Johnston, K. V., & Hernquist, L. 2005, *ApJ*, **632**, 872
- Sanders, J. L., & Das, P. 2018, *MNRAS*, **481**, 4093
- Sanderson, R. E., Wetzel, A., Loebman, S., et al. 2020, *ApJS*, **246**, 6
- Schlaufman, K. C., Rockosi, C. M., Allende Prieto, C., et al. 2009, *ApJ*, **703**, 2177
- Schlaufman, K. C., Rockosi, C. M., Lee, Y. S., et al. 2012, *ApJ*, **749**, 77
- Schlaufman, K. C., Rockosi, C. M., Lee, Y. S., Beers, T. C., & Allende Prieto, C. 2011, *ApJ*, **734**, 49
- Sharma, S., Bland-Hawthorn, J., Johnston, K. V., & Binney, J. 2011, *ApJ*, **730**, 3
- Smith, M. C., Evans, N. W., Belokurov, V., et al. 2009, *MNRAS*, **399**, 1223
- Steinmetz, M., Zwitter, T., Siebert, A., et al. 2006, *AJ*, **132**, 1645
- van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *CSE*, **13**, 22
- Veljanoski, J., Helmi, A., Breddels, M., & Posti, L. 2019, *A&A*, **621**, A13
- Vincenzo, F., Spitoni, E., Calura, F., et al. 2019, *MNRAS*, **487**, L47
- Walker, I. R., Mihos, J. C., & Hernquist, L. 1996, *ApJ*, **460**, 121
- Wetzel, A. R., Hopkins, P. F., Kim, J.-h., et al. 2016, *ApJL*, **827**, L23
- White, S. D. M., & Rees, M. J. 1978, *MNRAS*, **183**, 341
- Yanny, B., Rockosi, C., Newberg, H. J., et al. 2009, *AJ*, **137**, 4377
- Yanny, B., Newberg, H. J., Kent, S., et al. 2000, *ApJ*, **540**, 825
- Yuksel, H., Kistler, M. D., & Stanev, T. 2009, *PhRvL*, **103**, 051101