OPEN ACCESS



The Emptiness Inside: Finding Gaps, Valleys, and Lacunae with Geometric Data Analysis

Gabriella Contardo D. David W. Hogg 1,2,3 D. Jason A. S. Hunt D. Joshua E. G. Peek 4,5 D., and Yen-Chi Chen Chen Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Avenue, New York, NY 10010, USA; gabriella.contardo@sissa.it Center for Cosmology and Particle Physics, Department of Physics, New York University, 726 Broadway, New York, NY 10003, USA Max-Planck-Institut fur Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA
Department of Physics & Astronomy, Johns Hopkins University, 3400 North Charles Street, Baltimore, MD 21218, USA
Department of Statistics, University of Washington, Seattle, WA 98195, USA

Received 2021 December 23; revised 2022 September 5; accepted 2022 September 24; published 2022 October 31

Abstract

Discoveries of gaps in data have been important in astrophysics. For example, there are kinematic gaps opened by resonances in dynamical systems, or exoplanets of a certain radius that are empirically rare. A *gap* in a data set is a kind of anomaly, but in an unusual sense: instead of being a single *outlier* data point, situated far from other data points, it is a region of the space, or a set of points, that is anomalous compared to its surroundings. Gaps are both interesting and hard to find and characterize, especially when they have nontrivial shapes. We present in this paper a statistic that can be used to estimate the (local) "gappiness" of a point in the data space. It uses the gradient and Hessian of the density estimate (and thus requires a twice-differentiable density estimator). This statistic can be computed at (almost) any point in the space and does not rely on optimization; it allows us to highlight underdense regions of any dimensionality and shape in a general and efficient way. We illustrate our method on the velocity distribution of nearby stars in the Milky Way disk plane, which exhibits gaps that could originate from different processes. Identifying and characterizing those gaps could help determine their origins. We provide in an appendix implementation notes and additional considerations for finding underdensities in data, using critical points and the properties of the Hessian of the density.⁷

Unified Astronomy Thesaurus concepts: Astronomy data analysis (1858); Computational astronomy (293); Astrostatistics techniques (1886); Milky Way dynamics (1051)

1. Introduction

Hypothesis generation is a fundamentally unsolved problem in astronomy, and even more so in the era of large data sets. We have powerful tools for testing existing hypotheses against the big data sets of the 2020s, but we do not have a clear path for wholly new, unanticipated discoveries when the data scales reach the petabyte regime and data cannot easily be inspected by eye. Some significant progress has been made in finding rare anomaly objects in astronomical data sets, including using citizen science approaches (e.g., Hanny's Voorverp object, Lintott et al. 2009; green pea galaxies, Izotov et al. 2011; Boyajian's star, Boyajian et al. 2016) and machine-learning approaches (Baron & Poznanski 2017; Margalef-Bentabol et al. 2020; Martínez-Galarza et al. 2021; Storey-Fisher et al. 2021), as well as simply pure serendipity, as pointed out in Harwit (2019). In this work, however, we focus on automated detection of the paucity of sources, in the form of gaps in the density distribution. Hence, instead of looking for anomalous, rare (that is, "outlying") objects, we are looking for regions of the space (or sets of objects) that have anomalously low density (that is, in comparison to their immediate surroundings).

Original content from this work may be used under the terms of the Creative Commons Attribution 4.0 licence. Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Gaps (local underdensities in the distribution) often have significance to our understanding of the universe, or represent important discoveries. For instance, the "green valley" (Strateva et al. 2001; Baldry et al. 2004, 2006; Schawinski et al. 2014) is a gap in the relationship between galaxy mass and specific star formation. In this case, it is thought that the gap is caused by fast evolution of galaxies across the gap from the star-forming sequence to the passively evolving quiescent galaxies. Another gap has been discovered in the Gaia DR2 color-magnitude diagram of stars in Jao et al. (2018). It is thought to represent either ³He instabilities or transitions to full convection in M dwarfs (Feiden et al. 2021). "Gaps" in stellar population have also been of interest, such as the rotational gap as defined by Barnes (2010), which was observed in Kepler data by McQuillan et al. (2013) and in K2 data by Gordon et al. (2021). In the exoplanet population, the "radius valley" has also been under scrutiny, and analysis of this "gap" in the radius of exoplanets as a function of other properties (such as orbital period, stellar mass, and stellar age; Fulton et al. 2017; Berger et al. 2020; Gupta & Schlichting 2020; David et al. 2021) has been conducted to further understand its underlying mechanisms and causes. And the gaps in Saturn's rings (Holberg et al. 1982) and the asteroid belt (Dermott & Murray 1983) reveal important dynamical resonances. Similar kinds of gaps appear in velocity space in the local parts of the Milky Way disk, although it is mainly the overdensities (ridges) that have been studied for now (e.g., Antoja et al. 2018; Kawata et al. 2018). Their properties could be indicative of their respective origins, thus shedding light on our Galaxy's structure and its evolution. We focus on this application in the

⁷ A Python implementation of t methods presented here is available at https://github.com/contardog/FindTheGap.

remainder of the paper to illustrate our methods. Finally, it is also important to understand nonphysical gaps in our data, as caused by instrumental imperfections or errors in analysis. While these gaps will not immediately reveal new truths about our universe, identifying them in an automated way will help accelerate science.

In our conception, a gap is a region of locally lower density inside a point cloud, such that as you move away from the gap (in some or most directions) the density rises. To be more specific, a point is in a gap if there is (at least) one straight line you can draw such that the density locally rises away from the point in both directions along that line segment. Thus, a gap can be linear or planar or hyperplanar, and it is not required, in our conception, to be spatially compact. Furthermore, we do not consider that the surrounding density, outside the gap, must be homogeneous, not even approximately. Additionally, our focus is not on finding the "emptiest" gaps (i.e., with lowest density) but the (locally) deepest or steepest gaps: We aim at building methods that can detect not only low-density gaps but also regions or subspaces of substantial density but that lie within even higher-density regions.

In that sense, our definition of a gap is in contrast to the usage of the word "void" as it is conventionally used in, for example, the study of large-scale structure in cosmology, where void statistics have been important (see, e.g., Lavaux & Wandelt 2010; Hamaus et al. 2016; Kreisch et al. 2022). Voids are defined to be compact (sometimes even spherical) regions of zero or near-zero galaxy density within the distribution of galaxies, which is homogeneous and isotropic on large scales. Voids are found with high completeness with tessellations or even counts of galaxies in spherical subvolumes (Neyrinck 2008; Sutter et al. 2015; Banerjee & Dalal 2016). These methods are appropriate when the goal is to find lower-than-mean density regions in a homogeneous (on large or intermediate scales) distribution. That is excellent for the distribution of galaxies in 3-space but suboptimal when the target is an arbitrary point cloud in an arbitrary data space, where no homogeneity can be assumed, on any scale.

Geometric data analysis and topological data analysis (TDA) have proposed tools to explore the properties of structure in data space, notably by relying on derivatives and second derivatives of a density estimate. We only briefly overview here some major topics of interest in TDA, but for a more thorough overview of TDA, we refer the reader to Wasserman (2018), and to Chen (2017) for a special focus on kernel density estimators in this context. Several works have explored the problem of mode-finding (i.e., finding the maxima within the data), proposing methods, for instance, using the mean-shift algorithm, to find those modes and then perform mode-clustering (Cheng 1995; Comaniciu et al. 2002; Li et al. 2007; Chacón 2015). Another concept of interest in TDA is the Morse-Smale (MS) complex (Morse 1925), which is a way to partition the data space using "critical points" (maxima and minima) and the density gradient (using "ascending flows" and "descending flows"). It has been used notably to visualize the multivariate density function in Chen et al. (2017). Critical points can provide starting points to find local underdensities in the distribution (see the Appendix for more details). However, while the MS complex could be used, in principle, to define gaps (e.g., looking at the gradient flows from local minima to nearby saddle points), it will become computationally challenging, as the number of dimensions increases, to cover the various dimensions in which gaps can lie, or will be restricted to 1D gaps.

TDA is also often associated with persistent homology, which studies how topological features change as a function of scale. While persistence diagrams provide information on the topological features of a set of points (e.g., number of loops or connected components), it does not provide "localized" information (e.g., where these features are). The problem of interest that is the closest to our work is the problem of ridge detection. It has been explored, for instance, in Genovese et al. (2014), whose approach is very adjacent to ours: their goal is to extend on mode-finding and uncover hidden structures, in the form of overdensities, in the data, using projections of the Hessian. However, their method focuses on overdensities and relies on optimization to find the ridges (i.e., it does not provide a "score" or a metric that can be computed at any given point). Their approach will not be adaptable to find underdensities of any shape (e.g., not 1D) and can become computationally expensive compared to ours.

Another approach related to our problem is a clustering method presented in Zhang et al. (2007): in order to identify the clusters in a set of points, the authors propose to find the valleys and use those to separate the groups. To do so, they rely on the normalized density derivative and an approximation of the local convexity of the density (similar to the approaches we propose), using nonparametric density estimation based on neighbor numbers. However, the method requires computing the pairwise distance between each point, which will rapidly get computationally challenging as the data set grows.

In this paper, we present a statistic that can be computed at (almost) any point in the data space and that can be used as an estimator of "gappiness" of a region. This statistic permits the highlighting of a wider variety of gaps (in terms of shapes and dimensions) in a more generic fashion than methods relying on critical points (or mode-finding) or "ridge-finding" approaches (which are, additionally, not designed originally for underdensities but for overdensities). To the best of our knowledge, there are no other methods that address the problem of gap detection and characterization in this form. Furthermore, as we propose a statistic that can be directly computed for a given point, the finding of gaps does not require optimization per se, contrary to other methods. This results in potential computational advantages. Interestingly, our statistic can be easily reversed to perform overdensity detection (thus including "ridges").

This paper is organized as follows: Section 2 presents the data set used as an example use case throughout this paper and motivates our gap characterization goal in this context. We define our statistic in Section 3 and showcase how it performs on real data. Section 4 provides a discussion of the results we observe on this application and the advantages and limitations of our current method.

Additionally, we provide in the Appendix additional considerations and implementation notes: we propose a specific twice-differentiable density estimator with finite support (which is used throughout this paper) that can alleviate some computational issues compared to, e.g., a classical Gaussian kernel. We also comment on other possible ways to identify and trace gaps, using notably critical points, and we provide possible methodologies to do so. These can provide different properties than the statistic presented in the main paper that can be relevant to specific use cases.

A Python implementation of the methods and the density estimator is available at https://github.com/contardog/FindTheGap.

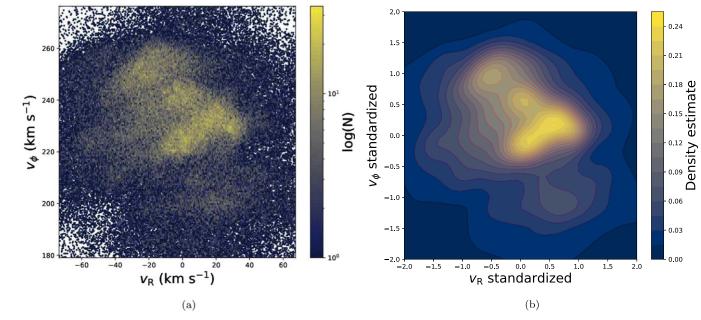


Figure 1. (a) 2D data set of velocities v_{ϕ} and v_{R} visualized as a hexagonal binning plot. (b) Density estimated with our quadratic kernel density estimator (see implementation notes in the Appendix), with a bandwidth $\Delta = 0.15$. Data are rescaled using standard normalization beforehand.

2. Example Data

We test and demonstrate the methods presented here on Galactic velocity and position data, where the identification and characterization of gaps are of crucial importance. We provide more details on the different data sets we build in the remainder of this section.

The second data release (DR2; Gaia Collaboration et al. 2018) from the European Space Agency's (ESA) Gaia mission (Gaia Collaboration et al. 2016) revolutionized our view of the Milky Way by providing position on the sky, parallaxes, and proper motions for over a billion stars across a large portion of the Galaxy. It also provided radial velocities for around 7 million stars, mostly within a few kiloparsecs of the solar neighborhood. This 6D phase-space sample revealed numerous disequilibrium features in the positions and kinematics of stars in the solar neighborhood and beyond.

Such disequilibrium features manifest as ridges and gaps in various dimensions. For example, Antoja et al. (2018) found a striking spiral pattern in the distribution of vertical position, z, versus vertical motion, v_z , which shows that the Milky Way is still phase mixing after some vertical perturbation, e.g., the passage of a satellite such as the Sagittarius dwarf galaxy (e.g., Antoja et al. 2018) or the buckling of the Galactic bar (Khoperskov et al. 2019). Antoja et al. (2018) and Kawata et al. (2018) also found ridges in the Galactocentric rotation velocity v_{ϕ} as a function of Galactic radius R, which can be signatures of the same satellite passage (e.g., Khanna et al. 2019; Laporte et al. 2019), Galactic spiral arms (e.g., Hunt et al. 2018), or resonances from the Galactic bar (e.g., Fragkoudi et al. 2019), or most likely a complex combination of all three.

These ridges and gaps are an extension of the long-known structure in the local v_R – v_ϕ kinematics across the observable disk. However, the change in the location of the ridges or gaps in kinematic space as a function of position in the Galaxy can shed light on their origin. For example, a gap with a resonant origin will move with a rate dependent on the order of the resonance in Galactic azimuth, ϕ . Developing methods that can not only find gaps in phase space but also quantify their rate of change in

higher-dimensional space will allow us to determine the origin of specific kinematic substructure, which in turn informs us on the structure and evolutionary history of our Galaxy. In this work we concentrate on the development and showcasing of the gapfinding algorithm at work in the Gaia data, and we defer the scientific interpretation of the substructure to future work.

For our sample, we use the recent intermediate data release (eDR3; Gaia Collaboration et al. 2021), which brings updated astrometry but no new radial velocity measurements, which are taken from DR2. We perform the photometric quality cuts that were suggested in Schönrich et al. (2019) for DR2, namely, we select stars with a color of $G_{\rm BP}-G_{\rm RP}<1.5$, a magnitude of G < 14.5, a fractional parallax error of $\pi/\sigma_{\pi} > 5$, a parallax uncertainty cut of σ_{π} < 0.1, a BP-RP excess flux factor of $1.172 < E_{\rm BPRP} < 1.3$, and with more than five visibility periods used, which may be overkill for eDR3. We derive distances naively as $d = 1/\pi$. Such an approximation is only valid for highly accurate parallaxes, yet the purpose of this work is to detect gaps, not make rigorous measurements. We use GALPY (Bovy 2015) to convert from the Gaia frame $(\alpha, \delta, \pi, \mu_{\alpha}, \mu_{\delta}, \nu_{R})$ to cylindrical coordinates $(R, \phi, z, v_R, v_\phi, v_z)$ assuming a distance to the Galactic center of $R_0 = 8.178$ kpc (Gravity Collaboration et al. 2019) and the Sun's height about the disk plane as 20.8 pc (Bennett & Bovy 2019). We calculate the vertical and azimuthal solar motion by combining R_0 with the proper-motion measurement of Sgr A* of $(\mu_b, \mu_b) = (-6.411 \pm 0.008, -0.219 \pm 0.007)$ (Reid & Brunthaler 2020). Thus, we have $v_{\odot} = 248.5 \text{ km s}^{-1}$ and $w_{\odot} = 248.5 \text{ km}$ 8.5 km s⁻¹. Finally, we perform an additional cut on the velocities, selecting stars with $-125 < v_R < 125$ and $100 < v_\phi < 300$.

From these data, we create the following data sets:

- 1. A 2D data set D_v , using v_R and v_{ϕ} , with an additional cut selecting stars within a distance of 200 pc from the Sun, resulting in a data set of \sim 217,000 stars. Visualization of this data set is provided in Figure 1(a).
- 2. Two 3D data sets, adding R and ϕ as a third dimension, respectively:

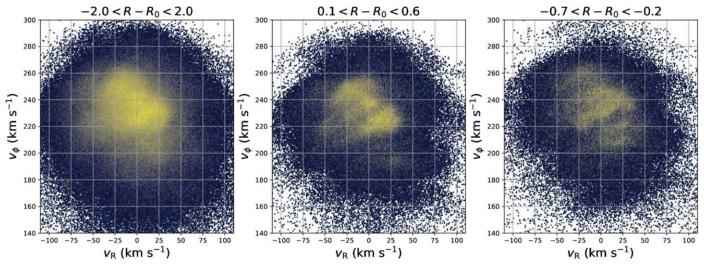


Figure 2. Projection of the 3D data set D_R (features v_{ϕ} , v_R , $R-R_0$) in velocity space, with all data (left panel), stars within $0.1 < R-R_0 < 0.6$ (middle panel), and stars within $-0.7 < R-R_0 < -0.2$ (right panel). Gaps in the distribution are most visible in the subsampled data (middle and right panels) and are visibly different for the two cuts, indicating a dependency with $R-R_0$.

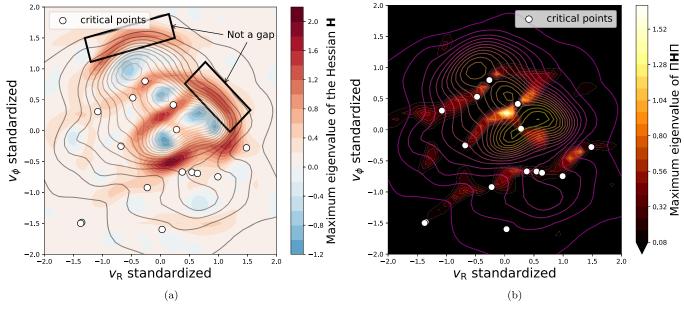


Figure 3. (a) Map of the maximum eigenvalue of the second derivative (Hessian) of the density estimated with a bandwidth $\Delta=0.15$ on a grid in the data space of D_{ν} , with density estimate shown as contour lines. We highlight a region with high values for λH that is, however, not a gap in the density distribution according to our definition, as it is not between two regions with higher density. White circles show the 15 critical points with the highest Hessian's maximum eigenvalues (see the Appendix for details). (b) Map of the maximum eigenvalue of $\Pi H\Pi$ for a bandwidth $\Delta=0.15$ on a grid in the data space of D_{ν} , with density estimate shown as contour lines. This criterion, compared to the maximum eigenvalue of the Hessian alone, efficiently removes areas with steep curvatures in the density that are, however, not gaps.

- (a) D_R using v_R , v_ϕ and $R-R_0$, with an additional cut selecting stars within 2 kpc in $R-R_0$ and 300 pc in ϕ , resulting in a data set of \sim 910,000 stars. Figure 2 shows the entire data set D_R in the velocity space, as well as two subsets of the data in subwindows in $R-R_0$. The gaps in those subsets are more visible, and one can see that their locations and shape (e.g., width) change from one subset to the other.
- (b) D_{ϕ} using v_R , v_{ϕ} and ϕ , with an additional cut selecting stars within 200 pc in $R R_0$ and within 15° in ϕ , resulting in a data set of \sim 752,000 stars.
- 3. A 4D data set $D_{R,\phi}$ including both velocities v_R , v_{ϕ} and both positions $R R_0$, ϕ . We keep stars within 2 kpc in R and within 15° in ϕ , resulting in a data set of \sim 1,670,000 stars.

We rescaled all data sets before computing density estimations and our methods throughout the paper, using standard normalization (other scaling methods could be used), in order to ensure coherent scale across dimensions. The methods presented here assume that all dimensions in the data are properly rescaled for the task at hand and/or according to underlying assumptions regarding the data. This might play a

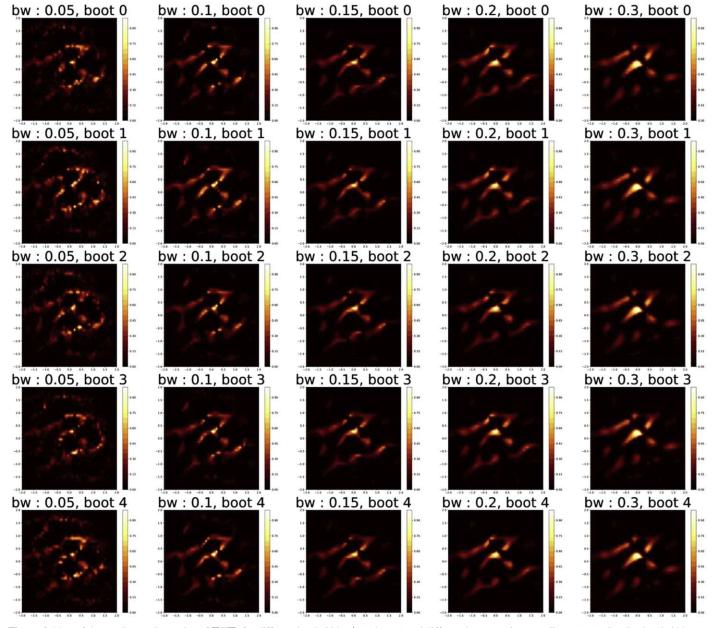


Figure 4. Map of the maximum eigenvalue of $\Pi H\Pi$ for different bandwidths Δ (columns) and different bootstrapping samplings (rows). Smaller bandwidths can grasp finer and smaller gap structures (e.g., bandwidth of 0.1); however, extremely low bandwidth becomes susceptible to finding spurious "gaps," which will not be stable or consistent across bootstraps (e.g., plots in the first column). At larger bandwidths, detected gaps become more stable across bootstraps, and the criterion can detect larger and wider structures.

crucial role in finding structures in the density distribution, as our kernel is symmetric in all dimensions.

3. Method

This section details the statistic we propose to estimate the "gappiness" of a point in the data space.

As a reminder, our goal is to provide a measure or statistic that can be interpreted as "how much a point can be considered to be lying within a gap," for any point in the space. Our definition of a gap is a region of *locally* lower density, such that there is (at least) one straight line along which density rises as we go away from the region in either direction along that line.

This definition is more generic than, for instance, relying on critical points, which are points where the gradient of the density is 0 (and could be minima, maxima, or saddle points; see the Appendix for more notes on that aspect). While a gap will have (at least) one critical point, its structure or surface can extend beyond that critical point. Our definition, without additional constraints, potentially extends a "gap" arbitrarily, as long as there is density rising on both sides.

Let us consider a data set (or point cloud) D of n points $x_i \in \mathbb{R}^d$, with boundaries (limits in \mathbb{R}^d) S. The approach we propose relies on a density estimator that can be of any nature, as long as it is twice differentiable. In the remainder of the paper, we consider a kernel-based density estimator (see the Appendix for the specific implementation used in our experiments), but the proposed statistic is agnostic to the actual estimator. We denote the density estimator p. As we consider a kernel-based estimator, p is associated with a bandwidth that we denote Δ . The estimator

 p_{Δ} is fitted on all the points within D. For any given point $x \in \mathbb{R}^d$, it is possible to compute the density estimate of that point and the gradient and second derivative (Hessian) of that density estimate. We denote the gradient vector of the density estimate ∇p and the Hessian matrix H, dropping the notation of the point of interest x for simplicity.

Following our definition of a gap, it makes sense to look at the properties of the second derivative (Hessian) of the density, as it is an indicator of the curvature of the density field. The maximum eigenvalue of the Hessian \boldsymbol{H} gives us an estimate of the regions of the space that have the biggest positive curvature in the density field: this is illustrated in Figure 3 on our 2D data set. Negative values correspond to "peaks" or "islands" in the density. High positive values align with regions where the density drops sharply (at that given bandwidth). These can be "valleys" (matching our gap definition), or they can be "cliffs": edges where the densities do not rise in any other direction (highlighted by black rectangles in Figure 3). These regions are not considered gaps in our definition.

Therefore, we need to examine the properties of the Hessian in a slightly different fashion, to account for this. The second derivative makes most sense to examine along directions where there is no first derivative, that is, where the second derivative delivers the most important nontrivial term in the Taylor series. Put another way, the second derivative confined to the subspace perpendicular to the gradient vector is the second derivative along all directions in which there is no first derivative. Because of that, every point in the d-dimensional data space can be thought of as a critical point (point of zero first derivative) in the (d-1)-dimensional subspace locally perpendicular to the local gradient at that point. Thus, it makes sense to examine the eigenvalues of the Hessian in that local subspace to ask whether the point is a local underdensity or local overdensity in the subspace. For our gappiness statistic, therefore, we propose to use the eigenvalues of the projection of the Hessian into the orthogonal subspace of the density gradient.

In that sense, our statistic relies on similar motivations to the ridge-finding optimization method presented by Genovese et al. (2014; low-dimensional overdensities in data). For $1 \le d' < d$ (where d is the number of dimensions of the data set), a d'-dimensional gap is characterized by a Hessian with d-d' large, positive values, and where the projection of the gradient on that subspace is 0.

We denote this projected second derivative tensor $\Pi H\Pi$, where H is the Hessian matrix and Π is a projection operator that projects into an orthogonal subspace, orthogonal to the gradient vector ∇p . If ∇p is seen as a column vector, then

$$\Pi \equiv \mathbb{I} - \frac{\nabla p \ \nabla p^{\top}}{\nabla p^{\top} \nabla p}, \tag{1}$$

where \mathbb{I} is the identity. This measure $\Pi H\Pi$ can be computed at any point in the data space. However, it will be undefined precisely at the true critical points (where the gradient vanishes), and indefinite for any truly empty gap (zero density at a given bandwidth).

We also note that other statistics, over $\Pi \boldsymbol{H} \Pi$ or its eigenvalues, could be interesting to explore, such as the trace of $\Pi \boldsymbol{H} \Pi$, which would provide different properties, especially at higher dimensions. Using the minimum eigenvalues of $\Pi \boldsymbol{H} \Pi$, on the other hand, will provide a "ridge" (or local overdensity) statistic.

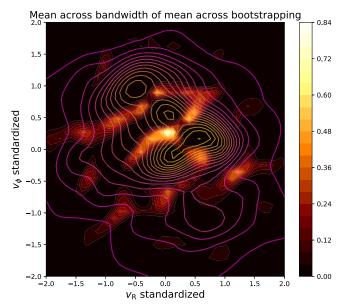


Figure 5. Map of the mean of the maximum eigenvalue of $\Pi H\Pi$ (rescaled between 0 and 1 for each run) across runs with different bandwidths and bootstrapping (see Figure 4). Taking summary statistics across bandwidths and bootstrapping allows us to highlight gaps of different sizes while keeping only the stable gap regions.

Figure 3(b) illustrates the map of the maximum eigenvalues of $\Pi H\Pi$ for all the points on a grid for the 2D data set D_{ν} , with the density estimate shown as contour lines, as well as the 15 critical points with the highest maximum eigenvalue of the Hessian (see the Appendix). We can see that this measure efficiently eliminates the irrelevant regions compared to using the maximum eigenvalue of the Hessian only and highlights the gap regions relevant to our definition. This visualization also shows that the regions with high values align with concave areas in the contour lines of the density, as expected given our target definition of a gap. We can see, however, that some of the "wider" gap regions (e.g., near 0.0 in ν_R and -1.0 in ν_ϕ) are not entirely covered and evaporate. This is related to the relationship between our gap statistics and the bandwidth Δ used by the density estimator.

Indeed, the choice of bandwidth will be critical to the properties of the gaps detected, as it will play a role, e.g., in the size of the gaps that can be detected. Deciding on a single bandwidth to do the analysis might be nontrivial and might leave out some important gaps. Small bandwidth might allow us to find smaller gaps, but they might be more sensitive to noise as well and will tend to undersmooth. This might make "spurious" gaps appear. On the other hand, large bandwidth will tend to oversmooth, and some gaps will disappear. There is also a relationship with the number of data points in the data set. The problem of bandwidth selection for kernel-based density estimators has been actively studied, and several methods have been proposed, such as Scott's rule of thumb (Scott 2015), Silverman's rule of thumb (Silverman 1998), and the Sheather and Jones method (Sheather & Jones 1991). among others. However, the problem remains unsolved in general. Therefore, ensuring the stability of our statistic evaluating the "gappiness" of a region/point in the data space is crucial. We propose to explore summary statistics of the maximum eigenvalues of $\Pi H\Pi$ when run across different bandwidths and using bootstrapping.

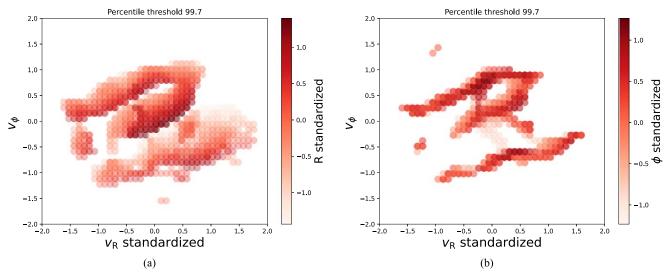


Figure 6. Regions in the 3D data sets (a) D_R and (b) D_{ϕ} with a maximum eigenvalue of $\Pi H\Pi$ above the 99.7th percentile on a $80 \times 80 \times 80$ grid, using a quadratic kernel density estimator with a bandwidth of 0.15. Selected regions are projected on the 2D velocity space and colored by their third dimension R and ϕ , respectively.

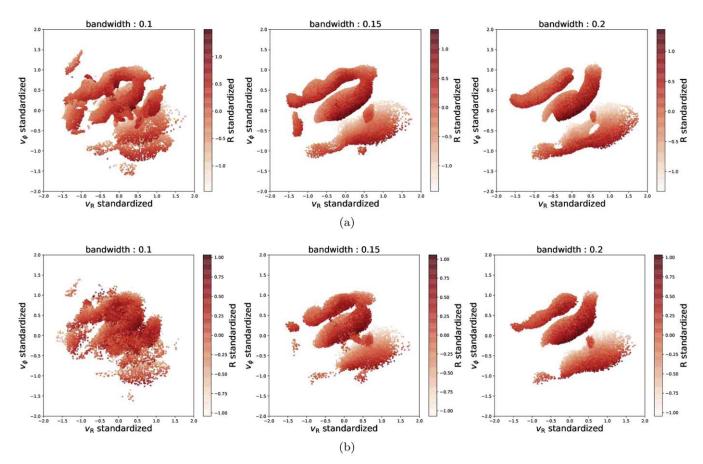


Figure 7. Data points selected in the (a) 3D data set D_R and (b) 4D data set $D_{R,\phi}$, such that their maximum eigenvalue of $\Pi H\Pi$ is above the 95th percentile and the 98th percentile, respectively. Points are colored by R.

Figure 4 shows the maximum eigenvalue of the $\Pi \boldsymbol{H} \Pi$ criterion for different bandwidths (columns) and different samplings (rows). It is apparent there that at the smallest bandwidth the gap estimation is dominated by noise, hence unstable per point. At larger bandwidth, some gaps disappear entirely, but the wider "valleys" get a better and more stable coverage.

By taking summary statistics across the different runs, it is possible to alleviate the instability of some gaps and to combine the stable gaps of different widths. First, we rescale each gap measurement map (maximum eigenvalue of $\Pi \boldsymbol{H}\Pi$) for each individual run so that its values lie between 0 and 1. Then, we take the mean across the different bootstrap run, in order to remove the "unstable" gaps at smaller bandwidths. Finally, computing the

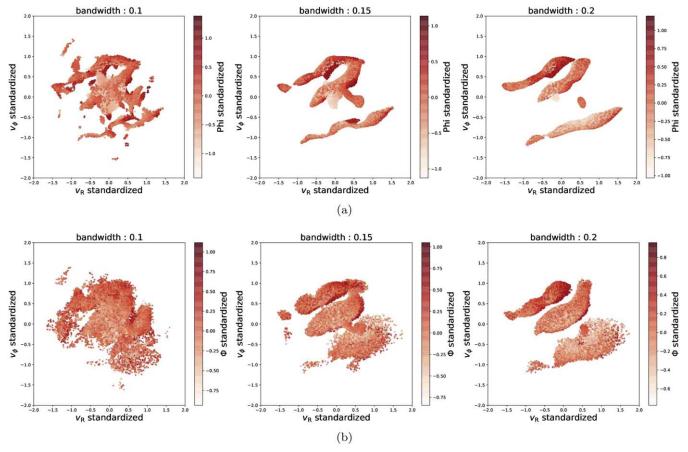


Figure 8. Data points selected in the (a) 3D data set D_{ϕ} and (b) 4D data set $D_{R,\phi}$, such that their maximum eigenvalue of $\Pi H\Pi$ is above the 95th percentile and the 98th percentile, respectively. Points are colored by ϕ .

mean of the mean maps across bandwidth allows us to keep the stable gaps of different width. We show in Figure 5 the final summary map averaging all the rescaled maps shown in Figure 4. Compared to Figure 3(b), using a summary across bandwidths allows us to better recover the wider gaps. However, we note that, in more general cases, users want to be careful in merging the statistic across bandwidths, as, in some instances, a bigger bandwidth might merge two smaller gaps together.

We also note that, although we explicitly perform bootstraps in Figures 4 and 5, it is probable that an explicit bootstrap is not required here to assess the shot noise (sample variance): the density estimate and its first and second derivatives are weighted sums over neighbors. Since we explicitly have a sum over weights to construct the density and its derivatives, we could look at the statistics of those weights to determine the shot noise variance in the density estimate. That would deliver further speed improvements to any pipeline making use of these tools.

Similarly, our "gap" statistic can be computed on our 3D data sets by gridding them and using the statistic as a selection threshold to highlight which regions are gaps in the distribution. Figure 6 shows the regions in D_R and D_ϕ with a maximum eigenvalue of $\Pi H\Pi$ above the 99.7th percentile, on a $80 \times 80 \times 80$ grid, with our quadratic kernel density estimator, with a bandwidth of 0.15. The selected regions are visualized in two dimensions v_R and v_ϕ , colored by their third dimension (R and ϕ , respectively). This clearly shows several gaps that move strongly as a function of R and more subtly in ϕ .

In many use cases and applications, it might be interesting to focus instead on which *data points* lie within gaps. This might prove useful from a scientific or data exploration point of view (studying the distribution and properties of the data lying in underdensities). It also provides an indirect way to trace the gap regions in a possibly less computationally expensive way than a fine-enough gridding would require, especially on higher-dimensional data sets. However, this "trick" will obviously fall short in the case of truly empty gaps.

Our gap statistic can be applied similarly to the data points within the data set directly. Depending on use cases, one can decide to keep all data points above a certain value of the maximum eigenvalue of $\Pi \boldsymbol{H} \Pi$ depending on its distribution (we note that this statistic is not a normalized statistics), to keep the data points above some percentile, or another criterion for cut.

We apply this selection process on our 3D data sets D_R and D_{ϕ} , integrating the positions in R and ϕ individually. Figures 7(a) and 8(a) show the data points above the 95th percentile for the maximum eigenvalue of $\Pi H\Pi$, for three different bandwidths, with the points colored by their respective additional position features R and ϕ . We see a very similar pattern to that in Figure 6, with gaps moving strongly as a function of R, and a more complex relationship with ϕ . We can apply a similar approach to our 4D data set $D_{R,\phi}$, combining both R and ϕ position. Figures 7(b) and 8(b) show the 98th percentile selection on our gap criterion computed with three different bandwidths (similar to the previous figures), respectively colored

by Φ and R as well. We can see that the "clusters" formed by the selected points become more blurry. While we show results with the same three bandwidths for the 3D and 4D cases for sake of illustration, it is unlikely that the "best" bandwidth for the 3D case will translate to the 4D case.

4. Results and Discussion

We present in this paper a statistic that can be used to highlight gaps, or local underdensities, in data distribution. We showcase the ability of our statistic to retrieve the observed gaps in the velocity distribution of nearby stars in the Milky Way and its potential to gain insights on the properties of those gaps. Our method relies on the use of a twice-differentiable density estimator. Such methods might themselves depend on a choice of bandwidth, which in turn impacts the possible gaps found by those methods. We illustrate that it is possible to combine different bandwidths sensibly to detect gaps of various widths, combined to a methodological way to ensure robustness of the gaps detected, using bootstrapping. When applied to our 3D and 4D data sets, we confirmed our expectations that the gaps in the kinematic space have dependencies with R and ϕ in terms of position and slope evolution. However, we defer further investigation on the underlying physics driving those gaps to future works.

Several paths for future investigation still remain. First and foremost, it remains nontrivial to go to much higher dimensions: one limitation will come from the natural limits of (kernel) density estimators to provide confident density estimates on higher-dimensional data sets. Besides this crucial problem, visualizing and extracting the characteristics of the gaps detected by our methods, when in higher dimensions, might become nontrivial as well.

Additionally, we presented here an analysis and tools using specific choices of "gap criterion" that approximate well our gap definition. But it would be interesting to investigate other possible criteria and statistics that could be computed on the gaps and underdensities. For instance, characterizing the "depth" of a gap instead might be relevant. However, such a measurement will be nontrivial to define and to compute efficiently. Another limitation of our methods might lie in the "summary" across bandwidths: in doing so, small gaps (in size) might merge and become difficult to distinguish. Exploring protocols to better handle those cases (e.g., in terms of bandwidth choice) will be crucial for some applications.

Conversely, similar investigations of "bumps" and ridges might also be of interest in many fields of applications. This problem has been more investigated in TDA, applied, for instance, on a cosmological data set to identify and characterize filaments of the cosmic web, as in, e.g., Xu et al. (2019). Our gap criterion as the maximum eigenvalue of $\Pi H\Pi$ can easily be reversed by taking the minimum eigenvalue instead, which will highlight the ridges and bumps in a density distribution.

It is a pleasure to thank Dan Foreman-Mackey (Flatiron) and Soledad Villar (Johns Hopkins University) for valuable discussions. Some of the ideas in this project were workshopped in the Astronomical Data Group Meetings at the Flatiron Institute and at the Machine Learning for Astronomy meeting at Max Planck Ringberg.

Software: This work used and benefited from the following Python libraries: astropy (Astropy Collaboration et al. 2018), galpy (Bovy 2015), jupyter (Pérez & Granger 2007), matplotlib

(Hunter 2007), numpy (Harris et al. 2020), pyTorch (Paszke et al. 2019), sklearn (Pedregosa et al. 2011), scipy (Jones et al. 2001).

Appendix Implementation Notes and Additional Considerations

In this appendix, we provide specific implementation notes and additional considerations on ways to find and characterize underdensities in data distribution that might be useful to readers in specific use cases.

We first describe a specific kernel for kernel density estimation that approximates a Gaussian kernel while providing finite support for faster computation. Then, we provide some comments on critical points and how they can be used to identify and trace specific types of gaps. We propose an algorithm for approximating the critical points and illustrate on our application how, using the gradient of the density estimate, one can trace "1D" valley gaps.

A.1. A Fast, Twice-differentiable Density Estimate

The statistic presented in this paper, as well as the methods presented below in this appendix, relies on the use of a twice-differentiable density estimator. The methods described are independent of the nature of the density estimator, so one could use whichever estimator that seems best for their applications, for instance, kernel density estimate, or even normalizing flows, as long as gradients and Hessians are available.

However, as a practical consideration, a Gaussian kernel, for instance, might become very slow for large data sets, as it is not compact. On the other hand, compact kernels, leading to faster kernel density estimators, are generally not twice differentiable at their edges. To address this issue, we propose to create a kernel that has finite support, is twice differentiable everywhere, and conforms to conventional ideas about bandwidth and resolution. It is the kernel used for all experiments shown in this paper.

We propose a custom quadratic kernel that approximates a Gaussian kernel near its center while offering finite support and being twice differentiable everywhere, including at the edge of the support. This density estimator is not properly normalized —it is designed to mimic a normalized Gaussian density estimator at the center of the kernel, rather than in the integral over the kernel. Nothing in our method relies on the density estimator being normalized, so this does not matter for our purposes, but if this density estimator is used for some integration or probability applications, it would make sense to normalize it correctly instead.

Namely, we build a kernel that is a close approximation to a Gaussian at small separations but goes smoothly to zero and has its slope and second derivative go smoothly to zero at what would be 3σ in the original Gaussian. This kernel k(s) can be expressed as a quadratic polynomial:

$$k(s) \equiv \begin{cases} 1 - 6\left(\frac{s}{3}\right)^2 + 8\left(\frac{s}{3}\right)^3 - 3\left(\frac{s}{3}\right)^4 & \text{for } s \leqslant 3\\ 0 & \text{for } s > 3 \end{cases}$$
 (A1)

$$s \equiv \sqrt{\frac{(x - x')^{\top}(x - x')}{\Delta^2}}$$
 (A2)

$$K(x - x') = \frac{1}{(2\pi \Delta^2)^{d/2}} k(s) , \qquad (A3)$$

where s is a dimensionless separation between a data point at d-dimensional location x and an evaluation point at d-dimensional location x' (computed with a Euclidean metric), and Δ is the bandwidth. Implicitly x and x' are $d \times 1$ column vectors. The factor of $(2\pi \Delta)^{d/2}$ normalizes the kernel comparably to the Gaussian it matches at small separations. The integral of the kernel over space is not unity.

Density estimation with this kernel can be made fast, because the kernel has compact support: in order to compute the density at a position x, it is only necessary to find neighboring data points out to separations of 3Δ in the x-space (or s=3 in the scaled separation variable s); we do not need all neighbors. Beyond separation s=3, the kernel and its first two derivatives exactly vanish everywhere. Density estimation with this kernel can be performed with only neighbors out to s=3. We obtain the list of $s \le 3$ neighbors exactly using the scipy (Jones et al. 2001) ckdtree implementation of the kd-tree, which has scales like $\ln N$ rather than N for problems like ours (where the bandwidth is smaller than the distribution of points). This can potentially speed up the density estimation enormously relative to methods that sum over all points.

A.2. 1D Gaps: Valleys in the Density Field

Critical points are points where the derivatives (slopes) of the density estimate are zero (i.e., they can be a minimum, a maximum, or a saddle). Since a gap will necessarily contain either a saddle point or a minimum point, critical points could be used to pinpoint gap regions in a local way.

First, we present an empirical way to approximate critical points in a given data set (or point cloud) in practice (although other methods could be used, e.g., as in Rosen et al. 2021).

Then, we propose a statistic that can be used to rank critical points in order to select the ones that lie in the gaps. We show on our application that this statistic is a sound proxy to select critical points lying in gaps in the distribution.

From these points, we then show how using the gradient of the density field can trace and highlight specific types of gaps, flowing from the saddle points, in the form of "valleys" or "streaks" in the density flow. This approach thus allows us to focus on "1D" gaps. We provide examples of our application to illustrate in which setups this might prove useful, e.g., for easier downstream analysis or visualization, by applying this "valley finder" on 2D slices from our 3D data sets.

A.2.1. Approximating Critical Points

To approximate the critical points within the boundary S of \mathbb{R}^d for a given estimator p_{Δ} , we propose the following methodology:

- 1. Create a grid of points G within the boundaries S, spaced by Δ , the bandwidth for the density estimator.
- 2. For each point $x \in G$, optimize to minimize the squared gradient of the density estimate. We use the BFGS algorithm provided in the *scipy* optimization library.
- 3. Group the resulting critical points so that points close to each other (e.g., within a distance function of the bandwidth Δ) are aggregated. In the following, we group together points that are closer than $0.01\Delta^2$ and summarize them as a single point being the mean of the group of points.

This process allows us to empirically approximate the set of relevant critical points (i.e., minimum, maximum, and saddles) in the density field, for a given bandwidth. However, we note that this method might be suboptimal in terms of computational efficiency, especially as the number of dimensions increases. We propose to use it here for the sake of simplicity and practicality for our use case. Other methods can be used, e.g., relying on contour trees as in Rosen et al. (2021).

A.2.2. Ranking Critical Points

We now need a way to rank the critical points in order of interest. First, it is relevant to point out that the nature of a critical point (i.e., if it is maximum, minimum, or saddle) can be determined through its index, i.e., the number of negative eigenvalues of the Hessian matrix \boldsymbol{H} (second derivative) of the density estimate at that point, since the eigenvalues of the Hessian give us indications on the direction of the curvature of the density field. A maximum point will have an index of d (number of dimensions of the data space), while a minimum point will have an index of 0. Critical points with indexes above 0 and below d are saddle points. This can potentially be used to reduce the set of critical points in the first place.

The statistic presented in this paper is undefined exactly at critical points and thus cannot be used as a "score" to rank the critical points. However, the maximum eigenvalue of the Hessian matrix is a good proxy criterion for our task, as it will rank higher the critical points that have the biggest positive curvature in the density field. While in theory this criterion selection could end up selecting critical points that are not corresponding to gaps by our definition (e.g., if the density field is in the form of a cliff with a perfectly flat bottom, a minimum critical point could land at the bottom there and have a high score), we observe that it is in practice a sensible choice, as the critical points "at cliffs" will move far enough away to have a lower maximum eigenvalue of the Hessian.

As an illustration, we compute the critical points estimated on the 2D data set D_{ν} described in Section 2, with a density estimator with a bandwidth Δ of 0.15. Figure 9(a) shows the resulting critical points colored by the value of the maximum eigenvalue of the Hessian at that location. We see that the critical points with the lowest value correspond to maximum, while the critical points with the highest values lie on saddle or local minimum regions. The critical points with values closer to zero tend to be in flatter, outer-skirt regions.

A.3. Tracing "1D" Valley Gaps

We now proceed to use the "best" candidate critical points (i.e., the ones with the highest scores) as a "starting point" to highlight gaps: indeed, while those points identify the local minimum point of the basin formed by a gap, they do not cover or characterize the entire gap "region" (or, in this instance, "valley") per se. However, we can see a saddle point as the "origin" of a valley gap, from which the gradient of the density can go either uphill (in the direction of the Hessian's eigenvectors with positive eigenvalues) toward higher-density regions or downhill (in the direction of the Hessian's eigenvectors with negative eigenvalues) toward lower-density regions. We propose to use these geometrical properties of the density field to trace the valley from a saddle point *x* in practice as follows:

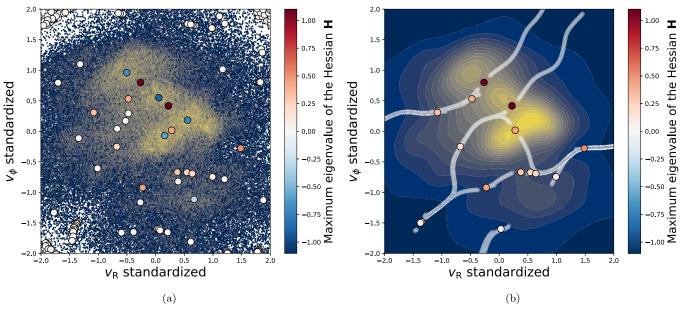


Figure 9. (a) Critical points found with a density estimator of bandwidth $\Delta = 0.15$, circled in black and colored by the value of the maximum eigenvalue of the Hessian of their density, for 2D data set D_R . Critical points with negative value (blue) correspond to maximum, while positive points correspond to local minimum and saddle. Points with values closest to zero lie mostly in the flat region of the density space. (b) Selection of N = 15 best critical points based on the value of the maximum eigenvalue of the Hessian of their density, with their "paths" (following the gradient direction) traced in white.

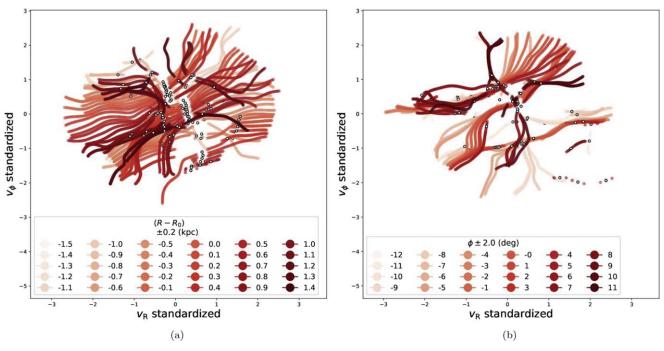


Figure 10. Best five critical points and their associated paths for each 2D "slice" in (a) D_R and (b) D_{ϕ} , colored by their respective R and ϕ window value, with density estimators of bandwidth 0.15.

- 1. Starting from a critical point x, take a small (e.g., a fraction of the bandwidth 0.1Δ) step away from x in the direction of the Hessian's eigenvector associated with the smallest Hessian's eigenvalue.
- 2. Descend "downhill" following the gradient of the density field in small steps of, e.g., a fraction of the bandwidth 0.1Δ until a low gradient is reached or until reaching the boundary S.

Figure 9(b) illustrates the paths we obtain with this process on the 2D data set D_{ν} , with a density estimator with a bandwidth

 $\Delta = 0.15$, selecting the best 15 critical points using the maximum eigenvalue of their Hessian. We can see that those paths neatly trace the visible valleys in the density distribution. Some of the selected critical points end up being connected to each other.

While tracing the "valleys" in this way restricts us to 1D gaps, it can provide an interesting method for a specific use case, if one has knowledge of the topology of the gaps of interest. Additionally, it might make the visualization and characterization of the properties of the gaps easier in higher

dimensions. We illustrate a different experimental protocol to analyze the gaps in a 3D data set, where we "slice" the original data set into 2D chunks. We then compute the critical points, rank and select the *K* "best" critical points, and finally compute their "paths" on slices (subset) of the original data instead, resuming to a 2D setting for each run.

Figure 10 illustrates this methodology on our original 3D data set D_R (Figure 10(a)) and D_{ϕ} (Figure 10(b)). The data sets were respectively sliced into 30 and 24 slices, creating 2D data sets with overlapping windows of similar size (0.2 kpc in $R - R_0$ for D_R , and 0.035 rad in ϕ for D_ϕ). The bandwidth used for the density estimator was 0.15. For each slice, we keep the best five critical points. This methodology highlights again the change in gap location as a function of R and ϕ . Additionally, we can visually see the critical points that seem to belong to the same gaps. The inclination and rate of change also look different for each gap. Another interesting aspect of this methodology, compared to the selection using our $\Pi H\Pi$ criterion in Figure 6, is that the "paths" can extend to regions that would not be selected using a $\Pi H\Pi$ -based cut (note the axis range of Figure 10). This approach might provide an easier path to properly analyze the rate and nature of the change, by characterizing each path individually, instead of the entire 3D structure or region selected through gridding.

We defer a deeper analysis of the characterization of the gaps and their links to physical origins to future works.

ORCID iDs

Gabriella Contardo https://orcid.org/0000-0002-3011-4784 David W. Hogg https://orcid.org/0000-0003-2866-9403 Jason A. S. Hunt https://orcid.org/0000-0001-8917-1532 Joshua E. G. Peek https://orcid.org/0000-0003-4797-7030 Yen-Chi Chen https://orcid.org/0000-0002-4485-306X

References

```
Antoja, T., Helmi, A., Romero-Gómez, M., et al. 2018, Natur, 561, 360
Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, AJ,
   156, 123
Baldry, I. K., Balogh, M. L., Bower, R., et al. 2006, MNRAS, 373, 469
Baldry, I. K., Glazebrook, K., Brinkmann, J., et al. 2004, ApJ, 600, 681
Banerjee, A., & Dalal, N. 2016, JCAP, 2016, 015
Barnes, S. A. 2010, ApJ, 722, 222
Baron, D., & Poznanski, D. 2017, MNRAS, 465, 4530
Bennett, M., & Bovy, J. 2019, MNRAS, 482, 1417
Berger, T. A., Huber, D., Gaidos, E., van Saders, J. L., & Weiss, L. M. 2020,
    J. 160, 108
Bovy, J. 2015, ApJS, 216, 29
Boyajian, T. S., LaCourse, D., Rappaport, S., et al. 2016, MNRAS, 457, 3988
Chacón, J. E. 2015, StaSc, 30, 518
Chen, Y.-C. 2017, Biostat. Epidemiol., 1, 161
Chen, Y.-C., Genovese, C. R., & Wasserman, L. 2017, EJSta, 10, 210
Cheng, Y. 1995, ITPAM, 17, 790
Comaniciu, D., Ramesh, V., & Del Bue, A. 2002, European Conf. on
   Computer Vision (Berlin: Springer), 561
David, T. J., Contardo, G., Sandoval, A., et al. 2021, AJ, 161, 265
Dermott, S. F., & Murray, C. D. 1983, Natur, 301, 201
```

```
Feiden, G. A., Skidmore, K., & Jao, W.-C. 2021, ApJ, 907, 53
Fragkoudi, F., Katz, D., Trick, W., et al. 2019, MNRAS, 488, 3324
Fulton, B. J., Petigura, E. A., Howard, A. W., et al. 2017, AJ, 154, 109
Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2018, A&A, 616, A1
Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2021, A&A, 649, A1
Gaia Collaboration, Prusti, T., de Bruijne, J. H. J., et al. 2016, A&A, 595, A1
Genovese, C. R., Perone-Pacifico, M., Verdinelli, I., & Wasserman, L. 2014,
    AnSta, 42, 1511
Gordon, T. A., Davenport, J. R., Angus, R., et al. 2021, ApJ, 913, 70
Gravity Collaboration, Abuter, R., Amorim, A., et al. 2019, A&A, 625, L10
Gupta, A., & Schlichting, H. E. 2020, MNRAS, 493, 792
Hamaus, N., Pisani, A., Sutter, P. M., et al. 2016, PhRvL, 117, 091302
Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Natur, 585, 357
Harwit, M. 2019, Cosmic Discovery: The Search, Scope, and Heritage of
   Astronomy (Cambridge: Cambridge Univ. Press)
Holberg, J., Forrester, W., & Lissauer, J. J. 1982, Natur, 297, 115
Hunt, J. A. S., Hong, J., Bovy, J., Kawata, D., & Grand, R. J. J. 2018,
   MNRAS, 481, 3794
Hunter, J. D. 2007, CSE, 9, 90
Izotov, Y. I., Guseva, N. G., & Thuan, T. X. 2011, ApJ, 728, 161
Jao, W.-C., Henry, T. J., Gies, D. R., & Hambly, N. C. 2018, ApJL, 861, L11
Jones, E., Oliphant, T., Peterson, P., et al. 2001, SciPy: Open Source Scientific
   Tools for Python, http://www.scipy.org/
Kawata, D., Baba, J., Ciucă, J., et al. 2018, MNRAS, 479, L108
Khanna, S., Sharma, S., Tepper-Garcia, T., et al. 2019, MNRAS, 489, 4962
Khoperskov, S., Di Matteo, P., Gerhard, O., et al. 2019, A&A, 622, L6
Kreisch, C. D., Pisani, A., Villaescusa-Navarro, F., et al. 2022, ApJ,
   935, 100
Laporte, C. F. P., Minchev, I., Johnston, K. V., & Gómez, F. A. 2019,
   MNRAS, 485, 3134
Lavaux, G., & Wandelt, B. D. 2010, MNRAS, 403, 1392
Li, J., Ray, S., & Lindsay, B. G. 2007, JMLR, 8, 1687, http://jmlr.org/papers/
   v8/li07a.html
Lintott, C. J., Schawinski, K., Keel, W., et al. 2009, MNRAS, 399, 129
Margalef-Bentabol, B., Huertas-Company, M., Charnock, T., et al. 2020,
    MNRAS, 496, 2346
Martínez-Galarza, J. R., Bianco, F. B., Crake, D., et al. 2021, MNRAS,
   508, 5734
McQuillan, A., Aigrain, S., & Mazeh, T. 2013, MNRAS, 432, 1203
Morse, M. 1925, AMS, 27, 345
Neyrinck, M. C. 2008, MNRAS, 386, 2101
Paszke, A., Gross, S., Massa, F., et al. 2019, in Advances in Neural Information
   Processing Systems 32, ed. H. Wallach et al. (Red Hook, NY: Curran
   Associates, Inc.), 8024
Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, JMLR, 12, 2825,
  http://jmlr.org/papers/v12/pedregosa11a.html
Pérez, F., & Granger, B. E. 2007, CSE, 9, 21
Reid, M. J., & Brunthaler, A. 2020, ApJ, 892, 39
Rosen, P., Seth, A., Mills, E., et al. 2021, Topological Methods in Data
   Analysis and Visualization VI (Berlin: Springer), 87
Schawinski, K., Urry, C. M., Simmons, B. D., et al. 2014, MNRAS, 440, 889
Schönrich, R., McMillan, P., & Eyer, L. 2019, MNRAS, 487, 3568
Scott, D. W. 2015, Multivariate Density Estimation: Theory, Practice, and
   Visualization (New York: Wiley)
Sheather, S. J., & Jones, M. C. 1991, J. R. Stat. Soc., B: Stat. Methodol.,
   53, 683
Silverman, B. W. 1998, Density Estimation for Statistics and Data Analysis
  (New York: Routledge), doi:10.1201/9781315140919
Storey-Fisher, K., Huertas-Company, M., Ramachandra, N., et al. 2021,
   MNRAS, 508, 2946
Strateva, I., Ivezić, Ž., Knapp, G. R., et al. 2001, AJ, 122, 1861
Sutter, P. M., Lavaux, G., Hamaus, N., et al. 2015, A&C, 9, 1
Wasserman, L. 2018, AnRSA, 5, 501
Xu, X., Cisewski-Kehe, J., Green, S. B., & Nagai, D. 2019, A&C, 27, 34
Zhang, C., Zhang, X., Zhang, M. Q., & Li, Y. 2007, PaReL, 28, 173
```