

MEASURING HUMAN ACTIVITY SPACES FROM GPS DATA WITH DENSITY RANKING AND SUMMARY CURVES

BY YEN-CHI CHEN^{*} AND ADRIAN DOBRA[†]

Department of Statistics, University of Washington, ^{}yenchi@uw.edu; [†]adobra@uw.edu*

Activity spaces are fundamental to the assessment of individuals' dynamic exposure to social and environmental risk factors associated with multiple spatial contexts that are visited during activities of daily living. In this paper we survey existing approaches for measuring the geometry, size and structure of activity spaces, based on GPS data, and explain their limitations. We propose addressing these shortcomings through a nonparametric approach called density ranking and also through three summary curves: the mass-volume curve, the Betti number curve and the persistence curve. We introduce a novel mixture model for human activity spaces and study its asymptotic properties. We prove that the kernel density estimator, which at the present time, is one of the most widespread methods for measuring activity spaces, is not a stable estimator of their structure. We illustrate the practical value of our methods with a simulation study and with a recently collected GPS dataset that comprises the locations visited by 10 individuals over a six months period.

1. Introduction. Collecting and statistical modeling of data on human movement in time and space is an important research endeavor in many fields, such as spatial epidemiology, demography and population science, urban design and planning, transportation research and environmental psychology (Apostolopoulos and Sonmez (2007), Chen et al. (2016), Dobra et al. (2017), Entwisle (2007), Hurvitz et al. (2014), Richardson et al. (2013)). Mapping individuals is difficult because a person's residence does not reflect their interaction with the physical and social environment (Kwan (2009)). Individuals spend considerable time away from their residences and traverse multiple administrative boundaries in their daily activities (Kwan (2013), Zenk et al. (2011)). For this reason, it is paramount to trace an individual through multiple spatial contexts to study environmental risk factors for disease (Cummins et al. (2007)). Statistical analyses that connect individuals to places by focusing on residential neighborhoods or administrative boundaries (e.g., census tracts) cannot capture short term but repetitive exposures to neighborhood-based risk factors (e.g., risk of violence or density of alcohol outlets). Going beyond the residential neighborhood of a person by collecting fine-grained positional data about where people actually spend time is especially relevant in studies that relate individual health to locally variable environmental factors (Basta, Richmond and Wiebe (2010)).

As human beings are inherently mobile, data about their spatiotemporal trajectories of travel are needed to construct relevant representations of their activity spaces. The notion of activity space has been introduced in the social sciences (Golledge and Stimson (1997)) and has its roots in the space-time-travel geography in which an individual's movements in time and space are conceptualized as space-time prisms (Hägerstrand (1963, 1970)). Activity spaces measure individual spatial behavior and capture individuals' experience of place in the course of their daily living through their observed location choices (Golledge (1999)).

Received July 2019; revised November 2019.

Key words and phrases. Activity space, global positioning systems (GPS), human mobility, kernel density estimation, space-time geography, topological data analysis.

They have been used to study the influence of the built environment on individuals' health-care accessibility (Sherman et al. (2005)). Activity spaces play a role in the study of social exclusion of individuals with low use of physical space which are less likely to be engaged in society (Schönfelder and Axhausen (2003)). Questions of interest relate to whether such individuals concentrate spatially or whether they are randomly scattered in the population? Are these individuals socially excluded from certain parts of the physical environment which could lead, for example, to lower chances of securing a job or higher costs of living? Activity spaces have also been used, among many applications, to assess segregation (Wong and Shaw (2011)), to measure exposure to food environments (Christian (2012), Kestens et al. (2010)) and to understand the geographic mobility patterns of older adults (Hirsch et al. (2014)).

Until about 15 years ago, research on activity spaces relied on locational data from travel diaries in which participants shared information about the trips they took in the past (Schönfelder and Axhausen (2003, 2004)). However, places outside the home neighborhood that are not socially significant are harder to be remembered and, consequently, they will be more likely to be missing from surveys. Smartphone-based location traces have recently become available for the study of human mobility and have proven particularly interesting, by providing the possibility of recording movements over time of individual people and aggregate movements of whole populations (Dobra, Williams and Eagle (2015), Williams et al. (2015)). This exciting new type of data holds immense promise for studying human behavior with precision and accuracy never before possible with surveys or other data collection techniques (Richardson et al. (2013)). Many high-resolution smartphone-based GPS location datasets have already been successfully collected and, subsequently, employed to assess human spatial behavior and spatiotemporal contextual exposures (Kwan (2012), Matthews and Yang (2013), Perchoux et al. (2013)), to characterize the relationship between geographic and contextual attributes of the environment (e.g., the built environment) and human energy balance (e.g., diet, weight, physical activity) (Berrigan et al. (2015), Zenk et al. (2011)), to study segregation, environmental exposure and accessibility in social science research (Kwan (2013)) or to understand the relationship between health-risk behavior in adolescents (e.g., substance abuse) and community disorder (Basta, Richmond and Wiebe (2010), Wiehe et al. (2013, 2008)). The wide array of completed and ongoing GPS studies provide key evidence that many people feel comfortable having their movements tracked (Zenk et al. (2012)).

In this paper we survey existing approaches for measuring the geometry, size and structure of activity spaces based on GPS data such as ellipses, shortest-path spanning trees and kernel density estimation, and explain the disadvantages of their use. To correct their shortcomings, we put forward a set of tools for measuring human activity spaces that comprise a nonparametric approach, called density ranking, and three types of summary curves. These curves fall within the broader domain of topological data analysis which is a flexible framework for detecting the structure and creating lower-dimensional summaries of distributions of complex or high-dimensional datasets (Carlsson (2009), Chazal and Michel (2017), Edelsbrunner and Harer (2008), Edelsbrunner and Harer (2010), Ghrist (2014), Kaczynski, Mischaikow and Mrozek (2004), Lum et al. (2013), Wasserman (2016, 2018)). The summary curves we discuss are based on level sets of density ranking which is closely related to level sets of a probability density function and the minimum volume set (Cadre, Pelletier and Pudlo (2013), Nuñez Garcia et al. (2003), Polonik (1997), Scott and Nowak (2006)).

The structure of the paper is as follows. In Section 2 we describe the GPS data we use to motivate and illustrate our developments. This is a never before analyzed dataset that comprises the spatiotemporal trajectories of daily living over a six months period of 10 individuals from a rural area. In Section 3 we present background on human activity spaces and describe existent methods for measuring them. In Section 4 we present density ranking, and in Section 5 we discuss three types of summary curves: the mass-volume curve, the Betti number

curve and the persistence curve. In Section 6 we introduce a novel mixture model for activity spaces and study its asymptotic properties. In Section 7 we present a simulation study. In Section 8 we apply density ranking and summary curves to the GPS data described in Section 2. Finally, in Section 9 we comment on the relevance of the proposed set of tools in the context of health research. We provide R scripts that implement our proposed methods at https://github.com/yenchic/density_ranking.

2. GPS data. We employ data from a GPS pilot study that involved three men and seven women that reside in a rural region. The study took place in 2016 with the approval of the local biomedical research ethics committee. These data have not been analyzed before. Each study participant was provided with a GPS-enabled Android smartphone for a period of six months. The smartphones together with their voice and data plans, whose costs have been covered by the pilot study, served as an effective incentive for study participation and adherence to the data collection protocol. The participants were asked to carry the smartphones with them at all times and also to keep them operational by regularly charging them. All 10 participants have been compliant with the protocol of the study and have returned their devices at the end of the study period.

The Android smartphones employ an assisted GPS system which produces accurate coordinate data with less battery power (allowing a phone to remain charged for at least 48 hours) than traditional GPS devices (e.g., GPS trackers). The positional data that were recorded contain timestamps, smartphone unique identifiers, latitude and longitude coordinates and information related to the accuracy of the reported coordinates (e.g., satellite connectivity). The smartphones were registered with a mobile device management (MDM) software that allowed the study personnel to manage, secure, monitor and track the smartphones from an easy to use online dashboard. The positional data were securely transmitted to a study database residing on a secure server over cellular or wireless networks using state of the art encryption techniques every time the smartphones had a data connection. The data were deleted from the smartphones immediately after transmission. This protocol guarantees that no confidential positional information could be accessed if a smartphone was lost or stolen.

The ages of the study participants were between 34 and 48 years. They share the same place of work. Their residences are located within a short commute of a couple of kilometers. The rural study area has a township in which most stores and markets are located. The local road network comprises a major primary road that traverses the township and several secondary roads. There are additional unpaved roads about which we did not have GIS data. The data comprise between 3500 and 8500 GPS locations for each of the 10 study participants. The MDM software installed on the phones was set to transmit a new location every time a device moved more than 250 meters. For this reason, more locations were recorded for those participants that traveled more. Figure 1 shows the GPS locations recorded for one of the study participants who was most active in the rectangular area shown in red in the left panel but also took several trips to more distant locations.

3. Existent approaches for measuring human activity spaces. Activity spaces represent the spatial areas within which an individual has direct contact during their daily travels. However, people do not move randomly in space. Due to the various preferences, needs, knowledge, constraints and limitations of movement, the areas visited by an individual are concentrated around one, two or more anchor locations that serve as origin and destination hubs of the routes followed by an individual. The anchor locations have key material or symbolic meaning for an individual: they include their home and work locations, together with, for example, the location of a child's school, favorite market and grocery store, a preferred entertainment venue or an airport. The probability of visiting a particular spatial location decreases as a function of its distance to the anchor locations and depends on its relative position

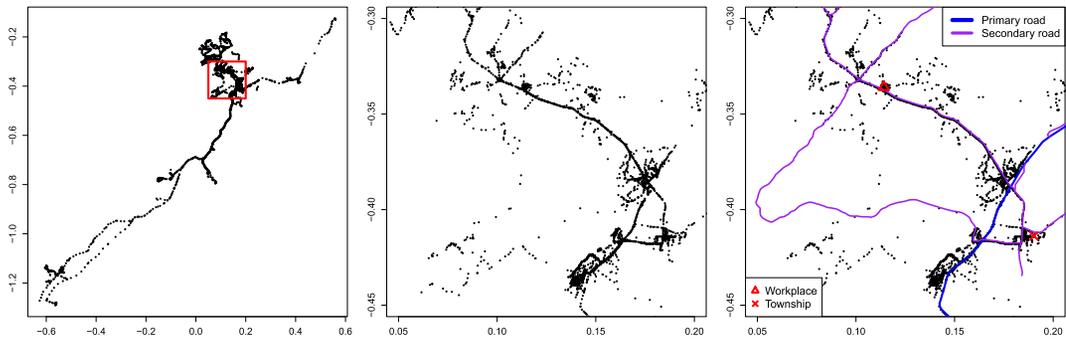


FIG. 1. Plots of the GPS data from one study participant. The left panel displays the complete GPS records of this individual. The latitude (x-axis) and longitude (y-axis) coordinates were shifted and scaled to preserve the privacy of the study participants. The middle panel represents the zoom-in area of the rectangular region shown in red in the left panel. In the right panel, the relevant GIS data was superimposed on the GPS locations: primary (blue) and secondary (purple) roads, the workplace (red triangle) of the study participants and the location of the township of the study area (red cross).

with respect to the most frequent directions of daily travel. The shape, structure and spatial extent of activity spaces are a function of the spatial configuration of the anchor locations and of the routes traveled between and around them (Schönfelder and Axhausen (2003)).

Characterizing the activity space of an individual involves: (i) determining the number and the spatial configuration of the anchor locations; (ii) identifying the places the individual is most likely to experience, in addition to the anchor locations, and differentiating these places from other places which the individual is considerably less likely to come in direct contact with; (iii) mapping the spatial configuration of these locations; and (iv) developing measures that quantify the geometry and spatial structure of the individual's activity space. Such measures capture the individual's degree of mobility while accounting for the underlying preferences for certain travel routes. Activity spaces are not designed to capture the maximal area in which an individual is active. Instead, they consist of one, two or more spatially contiguous areas structured around the anchor location in which an individual regularly engages in activities of daily living, together with the routes used by the individual to travel between these areas.

We denote by $\mathcal{T} = \{T_1, T_2, \dots, T_n\}$ the GPS positional data of an individual. The j th location is $T_j = (x_j, y_j, t_j)$ where x_j denotes latitude, y_j denotes longitude and t_j denotes the time when the location (x_j, y_j) was visited. We assume that $t_1 < t_2 < \dots < t_n$. The set of visited locations $\mathcal{X} = \{X_1, X_2, \dots, X_n\}$, where $X_i = (x_i, y_i)$, is the projection of \mathcal{T} onto the latitude and longitude coordinates. The times when the locations were visited together with the order in which locations were visited are lost through this projection. Information about the routes traveled by an individual are comprised in \mathcal{T} but are absent in \mathcal{X} . The set of anchor locations are denoted by $\mathcal{A} = \{A_1, A_2, \dots, A_{n_0}\}$. We note that \mathcal{A} is not necessarily a subset of \mathcal{X} since some anchor locations might need to be inferred from possibly noisy GPS measurements.

The existent literature has introduced several approaches for characterizing activity spaces. We describe them below, together with their advantages and limitations.

3.1. Ellipses. This appears to be one of the earliest and most popular methods for measuring activity spaces which has been concurrently developed in several research domains such as biological habitat research, transportation research and human geography (Schönfelder and Axhausen (2003, 2004)). Ellipses are fit to the set of visited locations \mathcal{X} based on knowledge of the most relevant anchor locations in \mathcal{A} such as residence and workplace.

There are two kinds of ellipses: the standard deviational, or confidence ellipse, and the home-work ellipse (Chaix et al. (2012)). The standard deviational ellipse is determined based on the assumption that the locations \mathcal{X} follow a bivariate normal distribution. This distribution can be centered around a central location determined as the arithmetic mean of the unique coordinates in \mathcal{X} or the weighted average by the frequency of visits at some locations. Since these averages might not designate an actual real-world address, the central location can be an anchor location, typically the home location which is recognized as the focal point of the lives of most people. The major axis of the standard deviational ellipse is the regression line of the latitude on the longitude coordinates, thus, the orientation of the ellipse reflects the sign of the correlation between coordinates. It is customary to report one and two standard deviational ellipses corresponding to 68% and 95% coverage probabilities (Sherman et al. (2005)).

The home-work ellipses differ from the standard deviational ellipses in that they are defined with respect to two anchor locations which become the two focal points of the ellipse. Typically, the focal points of the ellipse are selected to be the home and work locations. This defines the major axis of the ellipse. Its minor axis is determined by selecting one additional visited location, which could be another anchor location, or the most distant location in \mathcal{X} from the two focal points (Newsome, Walcott and Smith (1998)).

Measures that describe an activity space represented through the space inside an ellipse are the area of the ellipse, which expresses the extent of the activity space, and the ratio of the length of the major and minor axes which represents the relative extent to which an individual deviates from its most frequently used route (e.g., home to work and back) (Newsome, Walcott and Smith (1998)).

One major disadvantage of representing activity spaces through ellipses are their relatively inflexible geometry: the spatial distribution of activity locations is constrained to the shape of the ellipse. Locations inside the ellipse are considered to be likely places of daily activities, while the locations outside the ellipse are viewed as unlikely travel locations. This is a problem because the actual shape of activity spaces could be quite different than that of an ellipse and could comprise nonoverlapping spatial regions. Moreover, ellipses could suggest larger activity spaces since they capture the underlying variability of locations and are not robust to outliers. In addition, an ellipse imposes a symmetry of the activity space around its center even if half of the area covered by the ellipse does not contain any locations in \mathcal{X} . To get around these issues, Schönfelder and Axhausen (2004) proposed using amalgamations of ellipses constructed around two or more anchor locations (e.g., one ellipse having home location as its center, and another ellipse having the work location as its center), while Rai et al. (2007) have shown how to fit three other curved geometrical shapes: the Cassini oval, the bean curve and the superellipse which comprises a circle and an ellipse. Selecting one of these shapes is based on particular assumptions about the form of the activity space: one, two, three or four clusters of locations with or without intermediate locations between them. Nevertheless, determining which (if any) of these assumptions is appropriate for a certain spatial pattern of locations \mathcal{X} cannot be done without performing a visual inspection which is problematic for applications that involve a large number of mobility profiles.

3.2. Minimum convex polygons. In this approach the activity space of an individual is defined as the area delimited by the smallest convex polygon that contains all the locations in \mathcal{X} . This method has been applied to study both animal and human activity spaces (Buliung and Kanaroglou (2006), Fan and Khatkhat (2008), Lee et al. (2016), Worton (1987)). Although the determination of minimum convex polygons is computationally straightforward, they cannot properly capture the shape of an individual's activity space which is typically irregular due to certain areas in the proximity of the locations in \mathcal{X} being very unlikely to be visited

(e.g., inaccessible or undesirable locations). As such, they identify activity spaces as being spatially larger than other approaches (Hirsch et al. (2014)). Other shortcomings of minimum convex polygons include: (i) the anchor locations \mathcal{A} and other most frequently visited locations are not represented or even identified; (ii) they imply that an individual is active in only one contiguous spatial area; and (iii) outlier locations in \mathcal{X} can significantly change the coverage and the shape of the resulting activity spaces. The spatial extent of minimum convex polygons is typically measured using their area and perimeter, while their shape is measured through their compactness (Harding, Patterson and Miranda-Moreno (2013), Manaugh and El-Geneidy (2012)). This is a measure of how circular a polygon is defined as the ratio between the area and the perimeter squared, multiplied by 4π . Its values range from near 1 (a polygon very close to a circle) to near 0 (an elongated polygon close to a line). The shape of ellipses can also be measured using their compactness scores.

3.3. Shortest-path spanning trees. This method employs a more realistic representation of human travel: individuals most often move via road networks instead of by apparition or “as crow flies” from one place to another. As opposed to the other three approaches, which employ only the locations \mathcal{X} , the shortest-path spanning trees are constructed with respect to a road network that spans the reference area and also with respect to the order in which the locations in \mathcal{X} were visited. The routes followed by an individual during their daily travels are approximated by projecting the locations in \mathcal{X} on the road network, then by connecting each pair of consecutive locations (seen as an origin-destination trip) by the shortest path on the road network between them (Schönfelder and Axhausen (2003, 2004)). Golledge (1999) argues that road networks affect the individuals’ perception and knowledge of places; therefore, activity spaces should be based on the paths followed by the travelers. As such, the activity space of an individual is represented by the spanning tree that covers the part of the network defined by the union of the shortest road network paths that connect consecutive visited locations. The spanning tree can be measured using its length or using the total area of buffers with a fixed length (e.g., 200 meters) around the road network segments. These buffers attempt to capture the space around the road network segments that might be known to an individual by walking around (Kim and Ulfarsson (2015)). Anchor locations and segments that are more intensely used on the road network can be determined based on the visitation frequencies.

An advantage of the shortest-path spanning trees is that this approach moves away from the assumption that individuals have a continuous knowledge about the space around and between the locations they visit; ellipses and minimum convex polygons are based on this assumption. Their shortcomings come from their dependence on the availability of road network data. Such data might not have been collected at all or have lower quality in rural areas or in low resource countries. Moreover, if the visited locations are recorded at larger time intervals, approximating the route followed by an individual by the shortest path between two consecutive locations might be crude: the individual might have traveled significantly more than the shortest path would indicate.

3.4. Kernel density estimation. This approach considers raster grid cells that partition a wider area that includes the set of visited locations \mathcal{X} which is seen as a point pattern. An activity surface over this wider area is generated by assigning a value to each cell in the raster based on the distances from the center of the cell to the locations \mathcal{X} (Buliung (2001), Kwan (2000)). The probability that the individual that visited the locations \mathcal{X} was also active in a particular cell is proportional to the value assigned to that cell. The kernel density estimator (KDE) is the sum of “bumps” centered at the locations \mathcal{X} . The estimate of the bivariate

density at grid point x (also referred to as the intensity at x) is given by (Silverman (1986))

$$(1) \quad \hat{p}(x) = \frac{1}{nh^2} \sum_{i=1}^n K\left(\frac{d_i(x)}{h}\right).$$

Here, $K(\cdot)$ is a kernel, h is the bandwidth or smoothing parameter and $d_i(x)$ is the distance between the grid point x and the i th visited location $X_i = (x_i, y_i) \in \mathcal{X}$. The most usual choice for $K(\cdot)$ is a radially symmetric unimodal probability density function such as the bivariate normal density. However, since the number of visited locations \mathcal{X} could be very large for some GPS studies, it is preferable to employ a kernel that does not require evaluating the value of the kernel at all points in \mathcal{X} for every grid point. For example, consider the quartic (biweight) kernel function (Silverman (1986))

$$K_2(z) = \begin{cases} \frac{1}{\Xi_2}(1 - z^2)^2 & \text{if } |z| < 1, \\ 0 & \text{otherwise,} \end{cases}$$

where $\Xi_2 = \int_{\|x\| \leq 1} (1 - \|x\|^2)^2 dx$ is the normalizing constant to ensure that $K_2(\cdot)$ is a density function. With this choice, the KDE from equation (1) becomes

$$(2) \quad \hat{p}(x) = \frac{1}{\Xi_2 h^2} \sum_{d_i(x) < h} \left(1 - \left(\frac{d_i(x)}{h}\right)^2\right)^2.$$

Thus, locations in \mathcal{X} outside a circle with radius h centered at x are dropped in the evaluation of $\hat{p}(x)$. The probabilities of visiting grid cells that are at larger distances from the most frequently visited areas will be smaller compared to the probabilities of visiting grid cells that are at smaller distances. The choice of bandwidth h is very important as larger bandwidths give more smoothing. However, for the KDE in equation (2), h also represents the maximum distance of spatial interaction between locations. Therefore, the choice of h for a particular application could reflect the understanding of proximity and neighborhood in daily travel for the area in which the location data was collected (Schönfelder and Axhausen (2003)).

In the KDE approach, the activity space of an individual comprises all the grid cells with an estimated probability (density) of visitation above a certain threshold $\tau_1 > 0$. The anchor locations can be identified as those grid cells with an estimated probability density of visitation above a second threshold $\tau_2 \in (\tau_1, \infty)$. Kernel density estimation can identify activity spaces of any shape and can also estimate the corresponding anchor locations which is something the ellipse and the minimum convex polygon methods cannot do. The shortest-path spanning trees rule out locations that are not on the road network on which they were defined. For this reason, the KDE approach seems to be the most flexible existent approach for activity space determination. Measuring the resulting activity spaces can be done by calculating the area covered by the grid cells included in them. It is possible to eliminate some of these areas, if they are known to be unfavorable to activities of daily living (e.g., heavy industrial and utility areas), thereby refining the shape of the activity spaces (Schönfelder and Axhausen (2003, 2004)).

4. Density ranking. Despite its flexibility in measuring activity spaces, kernel density estimation sometimes fails to yield adequate results when applied to GPS datasets. Consider the left panel of Figure 2 in which we show the KDE of locations from the region in the middle panel of Figure 1. Although it correctly identifies two peaks with the highest concentration of locations, the KDE does not capture much of the underlying structure of the GPS data. In this section we discuss an alternative to KDE, called density ranking, that captures much more of the underlying mobility patterns of this individual; see the middle and right panels

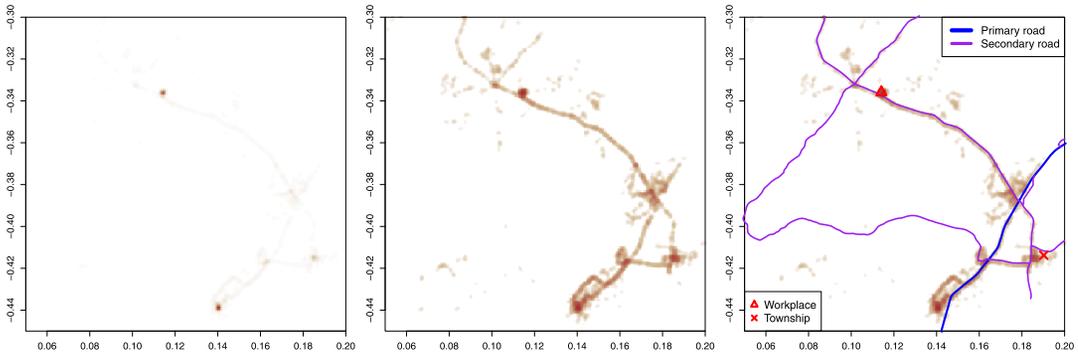


FIG. 2. A comparison between KDE and density ranking. In the left panel, we display the density contours from the KDE associated with the locations shown in the middle and right panels of Figure 1. In the middle panel, we show the contours identified by density ranking. In the right panel, we superimpose GIS data to the density ranking contours.

of Figure 2. It is apparent that many finer structures are not discernible using KDE, but they can be easily recognized when using density ranking. The KDE map only shows two grid cells that have high intensity: the workplace and another location that might be the home of this individual. On the other hand, the density ranking maps show the existence of numerous other grid cells located on the spatial trajectory followed by this individual. These regions represent the location of the township, road intersections or road segments.

Density ranking is a quantity derived from the KDE defined as

$$\hat{\alpha}(x) = \frac{1}{n} \sum_{i=1}^n \mathbb{I}(\hat{p}(X_i) \leq \hat{p}(x)),$$

where $\mathbb{I}(\cdot)$ is the indicator function. The density ranking function $\hat{\alpha}(x)$ is the fraction of observations in $\mathcal{X} = \{X_1, X_2, \dots, X_n\}$ whose estimated density is lower than the estimated density of the given point x . This function was called the α -function in Chen (2019). The density ranking function $\hat{\alpha}(x)$ is a probability-like quantity that takes values between 0 and 1. It has a natural relationship with the rank of the data points with respect to the KDE \hat{p} . Let $R_i = \sum_{j=1}^n \mathbb{I}(\hat{p}(X_j) \leq \hat{p}(X_i))$ be the rank of X_i with respect to \hat{p} . We have $R_i = 1$ if X_i has the lowest density, and $R_i = n$ if X_i has the highest density. Then,

$$\hat{\alpha}(X_i) = \frac{R_i}{n}$$

which implies that the density ranking at each observed data point is just the relative ranking of that point.

Density ranking has a straightforward interpretation related to the locations visited by an individual: for a point x with $\hat{\alpha}(x) = 0.8$, the probability density (measured by the KDE \hat{p}) at point x is higher than the probability density of 80% of all observed GPS locations. We say that x is in the region of the *top 20%* activity. Given a level $\gamma \in [0, 1]$, the level set of density ranking

$$\hat{A}_\gamma = \{x : \hat{\alpha}(x) \geq 1 - \gamma\},$$

can be interpreted as *the area of the top $\gamma \times 100\%$ activities*. The set \hat{A}_γ is the region within the contours of level $1 - \gamma$. Note that \hat{A}_γ is related to the *minimum volume set* (Nuñez Garcia et al. (2003), Polonik (1997), Scott and Nowak (2006)) and can be interpreted as a density level set with a probability content of γ (Cadre, Pelletier and Pudlo (2013)). In this view, \hat{A}_γ can be interpreted as an estimator of the smallest (in terms of volume) area covering at least $\gamma \times 100\%$ activities.

As explained in Section 3.4, given two prespecified levels τ_1, τ_2 with $\tau_1 < \tau_2$, the activity space based on the KDE \hat{p} comprises the grid cells x with $\hat{p}(x) \geq \tau_1$. The anchor locations are those grid cells x with $\hat{p}(x) \geq \tau_2$. Similar definitions of activity spaces and anchor locations can be given based on density ranking. We choose two levels $\gamma_1, \gamma_2 \in (0, 1)$ with $\gamma_2 < \gamma_1$. We define \hat{A}_{γ_1} to be the *top $\gamma_1 \times 100\%$ activity space or γ_1 -activity space*. Then, anchor locations are defined as the γ_2 -activity space. Choosing particular levels for the determination of human activity spaces or anchor locations can be done by examining the top 90, 80, . . . , 10% activity spaces. In Section 7 we show that the summary curves we introduce in the next section can be used to guide the choice of levels.

5. Summary curves. Based on density ranking, we obtain a two-dimensional function (a map) of human activity spaces. However, comparing maps associated with the activity patterns of multiple individuals is not straightforward without adequate summaries of the shape of these functions. To define such summaries, we use tools from topological data analysis (Chazal and Michel (2017), Edelsbrunner and Harer (2008), Wasserman (2016)). Specifically, we describe three types of summary curves that quantify the shape of a two dimensional function. These curves provide additional information about the geometry, size and structure of human activity spaces.

5.1. *Mass-volume curves.* Given a level γ , the size of the region \hat{A}_γ can be used to quantify an individual’s mobility in terms of the spatial extent of the γ -activity space. We measure size with the mass-volume function (Cl emen on and Jakubowicz (2013), Cl emen on and Thomas (2018), Nu n ez Garcia et al. (2003)) which is defined as

$$\hat{V}(\gamma) = \text{Vol}(\hat{A}_\gamma),$$

where $\text{Vol}(A) = \int_A dx$ is the volume of the set A . For example, if an individual has $\hat{V}(0.2) = 3 \text{ km}^2$, we say that the top 20% activities of this individual occur within a region of size 3 km^2 .

We subsequently define the *mass-volume curve* $\hat{V} = \{(\gamma, \hat{V}(\gamma)) : \gamma \in [0, 1]\}$ which describes how the volume of the γ -activity space \hat{A}_γ evolves when we vary the level γ . Mass-volume curves can be used to compare the degree of mobility of two individuals. Consider two example individuals with mass-volume curves \hat{V}_1 and \hat{V}_2 such that $\hat{V}_1(\gamma_0) > \hat{V}_2(\gamma_0)$ for some level $\gamma_0 \in [0, 1]$. We say that first individual has a higher mobility than the second individual in terms of the top $\gamma_0 \times 100\%$ activities.

5.2. *Betti number curves.* The mass-volume curve quantifies the activity space in terms of its size, but it does not provide any information about the shape of the activity space. Key concepts from topological data analysis (Chazal and Michel (2017), Edelsbrunner and Harer (2008), Wasserman (2016)) turn out to be very useful for this purpose. Two points in a set S are connected if and only if there exists a curve inside S that connects them. The set S is connected if any two points in the set are connected. The connected components of S are the induced partition from this relation. The connected components of S are a partitioning of S into subregions. Each connected component must not overlap with other connected components and must be connected. Two points that belong to two different connected components of S cannot be connected with a curve inside S .

We consider the connected components of the γ -activity space \hat{A}_γ . We define the Betti number function

$$\hat{\beta}(\gamma) = \text{number of connected components of } \hat{A}_\gamma$$

and the Betti number curve $\hat{\beta} = \{(\gamma, \hat{\beta}(\gamma)) : \gamma \in [0, 1]\}$. This curve captures how the number of connected components of the γ -activity space changes with the level γ . Note that the Betti

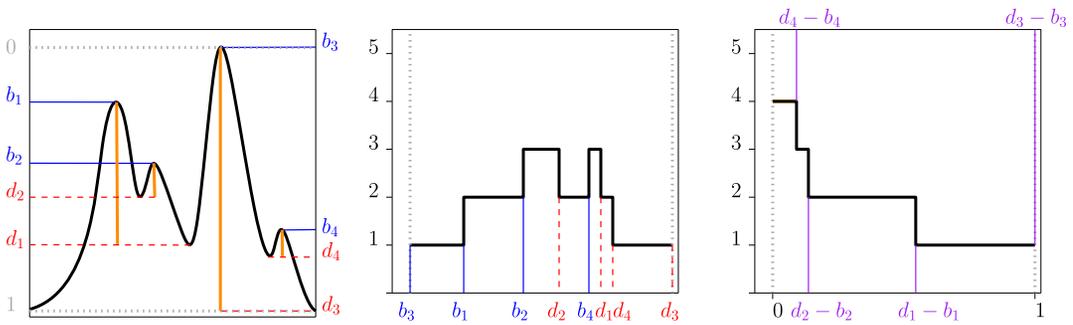


FIG. 3. An illustration of how the Betti number curve (middle panel) and the persistence curve (right panel) are computed for the function shown in the left panel.

number function is related to the number of local maxima of $\hat{\alpha}(x)$ and $\hat{\rho}(x)$. In the cluster analysis literature, the value $\hat{\beta}(\gamma)$ is interpreted as the number of clusters (Hartigan (1975), Rinaldo and Wasserman (2010)). The Betti number curve is closely related to the barcode plot in topological data analysis (Ghrist (2008), Wasserman (2016)).

We provide an example in Figure 3. The left panel shows a univariate function with four local maxima that correspond to levels b_1, b_2, b_3, b_4 , and four local minima that correspond to levels d_1, d_2, d_3, d_4 . A level set \hat{A}_γ comprises those regions where the function has a value above $1 - \gamma$. A new connected component is created when, as γ increases from 0 to 1, it passes one of b_1, b_2, b_3, b_4 . An existing connected component disappears when γ passes one of d_1, d_2, d_3, d_4 . Each of the orange vertical line segments represents a connected component: its upper and lower ends correspond to the birth and death time of this connected component, respectively. The middle panel shows the Betti number curve that corresponds to the function in the left panel. The Betti number curve goes up when the level γ hits the density ranking value of a local maximum and may drop when passing through the density ranking value of a local minimum or a saddle point. In this example, the Betti number curve increases whenever it passes b_1, b_2, b_3, b_4 and decreases when it passes d_1, d_2, d_3, d_4 . For a given value γ , the Betti number function $\hat{\beta}(\gamma)$ tells us the number of connected components in the top $\gamma \times 100\%$ activity region \hat{A}_γ . For example, if $\hat{\beta}(0.2) = 2$, the region of top 20% activities has two disjoint components. This implies that the individual's top 20% activities are concentrated around two areas that could correspond with the locations of this person's home and workplace. Individuals that record higher values of the Betti number function are those who tend to repeatedly visit a larger number of spatially distinct locations.

5.3. Persistence curves. The previous two types of curves focus refer to specific γ -activity spaces. Next, we define a third type of curve called a persistence curve that simultaneously considers all levels $\gamma \in [0, 1]$. The concept of persistence plays a key role in persistent homology, a branch of topological data analysis (Chazal and Michel (2017), Edelsbrunner and Harer (2008), Wasserman (2016)). The persistence curve is related to the accumulated persistence function (Biscio and Møller (2019)). The Betti number curve and the persistence curve can be viewed as functional summaries of topological features (Berry et al. (2018)).

We first define the persistence of a connected component. When we vary the level γ , new connected components may be created and existing connected components may disappear by merging with other connected components. We define the birth time of a connected component to be the level when this component is created and its death time to be the level at which this component disappears. A connected component is created at the level of the density ranking of a local mode and is often eliminated at the level of the density ranking of a local minimum or a saddle point. When two connected components merge into one, we

apply the following seniority rule (Chazal and Michel (2017), Wasserman (2016)): the older one (created at a lower level) stays alive while the younger one (created at a higher level) is eliminated. We define the death time of the connected components at level $\gamma = 0$ to be 0.

In the left panel of Figure 3, the two end points of an orange line segment correspond to the birth (creation) and death (elimination) of a connected component. The corresponding levels of the end points, b_ℓ and d_ℓ , are the birth time and death time of that connected component. There is a direct relationship between birth and death times and the Betti number curve: the Betti number increases by 1 whenever it passes the birth time of a connected component, and it decreases by 1 when it passes the death time of a connected component. In Figure 3, two connected components are created at levels b_1 and b_2 , and are eliminated at levels d_1 and d_2 . At b_1 and b_2 , the Betti number increases by 1, and it decreases by 1 at d_1 and d_2 . Since $b_1 < b_2$, the connected component created at b_1 is older than the connected component created at b_2 . When the two connected components merge at level $d_2 \in (b_2, d_1)$, the connected component created at b_1 remains, while the connected component created at b_2 is eliminated.

For each connected component, its persistence (also called life time) is the difference between the birth and the death time. In Figure 3 the length of an orange line segment is the persistence of that connected component. We define the persistence function

$$\widehat{\rho}(t) = \text{number of connected components whose persistence} \geq t$$

and the persistence curve $\widehat{\rho} = \{(t, \widehat{\rho}(t)) : t \in [0, 1]\}$. An example persistence curve is shown in the right panel of Figure 3. The persistence curve is a nonincreasing curve since we are thresholding on the life time of connected components. There will always be a connected component with a life time close to 1 because, by definition, the data point with the highest KDE value has rank equal to the sample size n , making its density ranking equal to 1. Due to the resolution of the underlying grid used, it is possible to see a connected component with life time close to but less than 1.

The persistence curve provides new information about the spatial distribution of the activity space. Unlike the mass-volume curve or Betti number curve that describes characteristics of level sets \widehat{A}_γ at particular levels γ , the persistence curve characterizes the collection of all level sets $\{\widehat{A}_\gamma : \gamma \in [0, 1]\}$. This is because, in order to compute the persistence of each connected component, we need to consider various levels to determine its persistence. An individual has a high-persistence curve when the corresponding density curve has many highly persistent connected components. These are regions this individual repeatedly visits: most likely, these represent their anchor locations. This type of information is not directly related to a particular γ -activity space. Instead, it is a quantity describing patterns across activity spaces at different levels.

6. A mixture model for human activity spaces. In this section we propose a statistical model that captures the most significant features of human activity spaces. We denote by P_{GPS} the probability distribution that defines the activity space of an individual. Note that the model we consider is a working model that provides insights on the data analysis; it may not reflect the actual data generating process. The observed locations $\mathcal{X} = \{X_1, \dots, X_n\}$ are assumed to be independent samples from P_{GPS} . We write this distribution as a mixture with three components

$$(3) \quad P_{\text{GPS}}(x) = \pi_0 P_0(x) + \pi_1 P_1(x) + \pi_2 P_2(x),$$

where $P_0(x)$ is an atomic distribution, $P_1(x)$ is a one-dimensional distribution, $P_2(x)$ is a two dimensional distribution and $\pi_0 + \pi_1 + \pi_2 = 1$ with $\pi_j \geq 0$ are proportions. The three components of the mixture (3) represent the key elements of the activity space represented by P_{GPS} : P_0 is a distribution that puts probability on the anchor locations \mathcal{A} ; P_1 is a distribution

describing the roads \mathcal{R} used by an individual when traveling between anchor locations; and P_2 is a distribution describing the areas \mathcal{O} around the anchor locations in which an individual moves. Although P_0 and P_1 do not have conventional probability density functions, they admit a generalized density function called the Hausdorff density (Mattila (1995), Preiss (1987)). Let $B(x, r)$ be the ball centered at x with a radius $r > 0$. For a positive integer s , the s -dimensional Hausdorff density (s -density) at x given P_{GPS} is

$$\mathcal{H}_s(x) = \lim_{r \rightarrow 0} \frac{P_{GPS}(B(x, r))}{C_s \cdot r^s},$$

where C_j is the volume of an j -dimensional unit ball ($C_0 = 1, C_1 = 2$ and $C_2 = \pi$). We denote by p_0, p_1 and p_2 the zero, one and two-dimensional Hausdorff densities given P_0, P_1 and P_2 . Namely, $p_0(x)$ is a mass at point x , $p_1(x)$ is a one-dimensional density value at x , and $p_2(x)$ is a two-dimensional density value at x . Furthermore, \mathcal{A}, \mathcal{R} and \mathcal{O} represent the support of p_0, p_1 and p_2 , respectively.

We define the dimension $\omega(x)$ of a point x with respect to P_{GPS} as follows:

$$(4) \quad \omega(x) = \begin{cases} 0 & \text{if } x \in \mathcal{A}, \\ 1 & \text{if } x \in \mathcal{R} \setminus \mathcal{A}, \\ 2 & \text{if } x \notin \mathcal{A} \cup \mathcal{R}. \end{cases}$$

The dimension of an anchor location is 0. The dimension of a location on a road inside the activity space that is not an anchor location is 1. The dimension of all the other locations is 2.

The following result shows that there are two equivalent ways to define the dimension $\omega(x)$:

THEOREM 1. *Given assumption (S) from the Supplementary Material (Chen and Dobra (2020)), Section 3, the definition of $\omega(x)$ in equation (4) is equivalent with the following two definitions:*

$$\omega(x) = \max\{s : \mathcal{H}_s(x) < \infty, s = 0, 1, 2\},$$

and

$$\omega(x) = \begin{cases} 0 & \text{if } p_0(x) > 0, \\ 1 & \text{if } p_0(x) = 0, p_1(x) > 0, \\ 2 & \text{if } p_0(x) = 0, p_1(x) = 0, p_2(x) \geq 0. \end{cases}$$

The proof of Theorem 1 is given in Supplementary Material (Chen and Dobra (2020)), Section 4. Using $\omega(x)$ and $p_j(x), j = 0, 1, 2$, we define a ranking comparison between two points x_1 and x_2 . We write $x_1 > x_2$ if

$$\omega(x_1) < \omega(x_2), \quad \text{or} \quad \omega(x_1) = \omega(x_2), \quad p_{\omega(x_1)}(x_1) > p_{\omega(x_2)}(x_2).$$

To compare two points, we first compare their dimension. The point that has a higher dimension is ranked higher. If both points have the same dimension, we compare their density value in that dimension. Then, the population quantity that density ranking is approximating is

$$(5) \quad \alpha(x) = P_{GPS}(x \geq X_1),$$

where X_1 is a random variable with distribution function P_{GPS} . Theorem 1 is needed to prove the next result that shows that $\hat{\alpha}(x)$ is a consistent estimator of $\alpha(x)$ which explains why density ranking yields stable results in measuring human activity spaces.

THEOREM 2. *Given assumptions (K1–2), (S) and (P1–2) from the Supplementary Material (Chen and Dobra (2020)), Section 3, and $\frac{nh^6}{\log n} \rightarrow \infty, h \rightarrow 0$, we have*

$$\int |\hat{\alpha}(x) - \alpha(x)|^2 dP_{\text{GPS}}(x) \xrightarrow{P} 0,$$

$$\int |\hat{\alpha}(x) - \alpha(x)|^2 dx \xrightarrow{P} 0.$$

The proof of Theorem 2 is given in the Supplementary Material (Chen and Dobra (2020)), Section 4.

The collection of anchor locations \mathcal{A} and roads connecting anchor points \mathcal{R} are of key interest and must be properly recovered from GPS data. Next, we show that the level set of $\hat{\alpha}(x)$, under suitable choices of level, will be a consistent estimator of \mathcal{A} and \mathcal{R} . Recall that $\hat{A}_\gamma = \{x : \hat{\alpha}(x) \geq 1 - \gamma\}$ is the level set of density ranking.

THEOREM 3. *Given assumptions (K1-2), (P) and (S0) from the Supplementary Material (Chen and Dobra (2020)), Section 3, and $\frac{nh^2}{\log n} \rightarrow \infty, h \rightarrow 0$, we have*

$$P_{\text{GPS}}(\hat{A}_{\pi_0} \Delta \mathcal{A}) \xrightarrow{P} 0,$$

where for sets A and B , $A \Delta B = (A \setminus B) \cup (B \setminus A)$ is their set difference and $P_{\text{GPS}}(A) = P(X_1 \in A)$, where X_1 has distribution P_{GPS} . Moreover, if we further assume (S1) from the Supplementary Material (Chen and Dobra (2020)), Section 3, we have

$$P_{\text{GPS}}(\hat{A}_{\pi_0 + \pi_1} \Delta (\mathcal{A} \cup \mathcal{R})) \xrightarrow{P} 0.$$

The proof of Theorem 3 is given in the Supplementary Material (Chen and Dobra (2020)), Section 4. The set difference Δ is a conventional measure of the difference between two sets. Applying the probability P_{GPS} to the set difference is a common measure of the convergence of a set estimator (Doss and Weng (2018), Mason and Polonik (2009), Qiao (2017), Rigollet and Vert (2009)). Theorem 3 shows that \hat{A}_{π_0} and $\hat{A}_{\pi_0 + \pi_1}$ are consistent estimators of \mathcal{A} and $\mathcal{A} \cup \mathcal{R}$, respectively. With this fact, we can use the difference $\hat{A}_{\pi_0 \pm \pi_1} \setminus \hat{A}_{\pi_0}$ as an estimator of \mathcal{R} . Namely, \hat{A}_{π_0} can be used to recover the anchor locations, and $\hat{A}_{\pi_0 + \pi_1} \setminus \hat{A}_{\pi_0}$ can be used to reconstruct the sections of the roads covered by an individual’s activity space.

Finally, we show that under the mixture model in equation (3), the KDE $\hat{p}(x)$ diverges with a probability tending to 1 at any anchor location or any point on a road connecting two anchor points.

THEOREM 4. *Under assumptions (K1–2) from the Supplementary Material (Chen and Dobra (2020)), Section 3 and $h \rightarrow 0$, we have $\mathbb{E}(\hat{p}(x)) \rightarrow \infty$ for any $x \in \mathcal{A} \cup \mathcal{R}$.*

The proof of Theorem 4 is given in the Supplementary Material (Chen and Dobra (2020)), Section 4. This result shows why the KDE does not give a stable estimator of \mathcal{A} and \mathcal{R} which explains why, for the activity space in Figure 2, the KDE does not properly detect its structure.

7. Simulation study. We consider an example individual whose activity space has three anchor locations \mathcal{A} : home (located at $(0, 0)$), office (located at $(0, 2)$) and gym (located at $(2, 0)$); see Figure 4. We assume that the anchor locations \mathcal{A} are connected by three straight segments of road \mathcal{R} . The individual spends 60% of their time in the anchor locations and 30% of their time traveling on the roads. In the rest of their time, this individual walks in the neighborhoods around their office and home, but never walks in the vicinity of their gym.

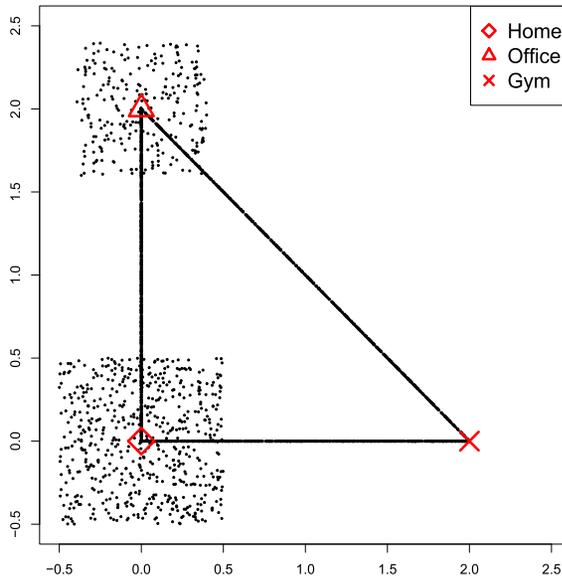


FIG. 4. Activity space for the simulation study, and scatter plot of $n = 8000$ locations sampled from the mixture model (6). The anchor locations are shown as follows: home (red diamond), office (red triangle) and gym (red cross).

When the individual is at an anchor location, they spend 50% of their time at home, 30% of their time at work, and 20% of their time at the gym. For this individual's activity space, the mixture model in equation (3) is written as ($\pi_0 = 0.6$, $\pi_1 = 0.3$, $\pi_2 = 0.1$):

$$(6) \quad P_{\text{GPS}}(x) = 0.6P_0(x) + 0.3P_1(x) + 0.1P_2(x),$$

with

$$P_0(x) = 0.5\delta_{(0,0)}(x) + 0.3\delta_{(0,2)} + 0.2\delta_{(2,0)}.$$

Here, $\delta_{(a,b)}(x)$ is a function that puts a point mass at (a, b) . The time in which the individual travels between the anchor locations is divided as follows: 30% on the road between home and gym, 20% on the road between gym and office and 50% on the road between home and office. We assume that the individual travels with the same speed on all road segments. For 70% of their total walk time, the individual moves uniformly within the square $[-0.5, 0.5] \times [-0.5, 0.5]$ centered at their home, and for remaining 30% the individual moves uniformly within the square $[1.6, 2.4] \times [-0.4, 0.4]$ centered at their office. With these assumptions, the distributions P_1 and P_2 in equation (6) are completely specified.

We generate $n = 8000$ samples from the mixture model (6); see Figure 4. We use a smoothing bandwidth of 0.5 to compute the density ranking. The corresponding contours (top left panel of Figure 5) show a very good agreement with the anchor locations and the road segments (top right panel of Figure 5). We determine two level sets of density ranking (see the bottom panels of Figure 5): $\hat{A}_{0.6}$ ($\pi_0 = 0.6$) and $\hat{A}_{0.9}$ ($\pi_0 + \pi_1 = 0.9$) corresponding with the mixture weights in Eq. (6). We see that $\hat{A}_{0.6}$ recovers all three anchor locations \mathcal{A} , while $\hat{A}_{0.9}$ recovers the anchor locations and the road segments $\mathcal{A} \cup \mathcal{R}$. This is consistent with our theoretical results, in particular, with Theorem 3.

We compare the relative performance of kernel density estimation and density ranking by simulating $n = 8000$ samples from the mixture model (6) 100 times. For density ranking, we determine the level sets $\{\hat{A}_\gamma : \gamma = 0.05, 0.10, \dots, 0.95\}$. For kernel density estimation, we determine the level sets $\{\hat{A}_{\gamma \cdot \max_x \hat{p}(x)} : \gamma = 0.05, 0.10, \dots, 0.95\}$. For each simulation experiment and each of level set A , we calculate the distances $P_{\text{GPS}}(A \Delta \mathcal{A})$ and $P_{\text{GPS}}(A \Delta (\mathcal{A} \cup \mathcal{R}))$.

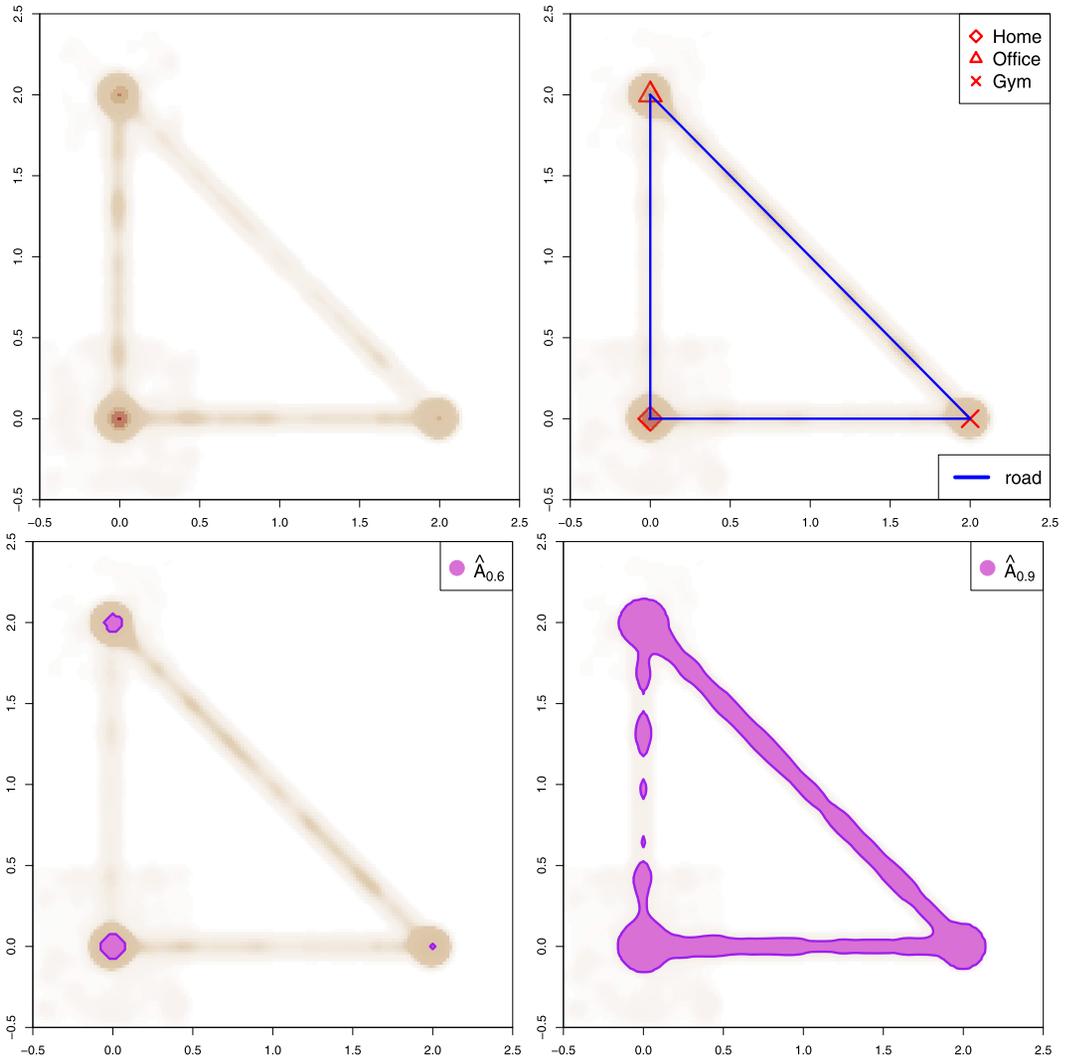


FIG. 5. Analysis of simulated data. Top left panel: Contours of density ranking. Top right panel: The anchor locations and the road segments superimposed on the density ranking contours. Bottom left panel: The level set $\hat{A}_{0.6}$ of density ranking. Bottom right panel: The level set $\hat{A}_{0.9}$ of density ranking.

These distances represent the error of estimating the anchor locations \mathcal{A} and the combined anchor locations and road segments $\mathcal{A} \cup \mathcal{R}$ with the level set A . The average estimation errors are displayed in Figure 6. The standard errors of the curves in Figure 6 are extremely small (0.003–0.005) and have been omitted.

For the purpose of estimating the anchor locations \mathcal{A} , the left panel of Figure 6 shows that the level sets from both kernel density estimation and density ranking work well, although the level sets from density ranking achieve a smaller error for levels below $\pi_0 = 0.6$ which represents the true percentage of time spent in the anchor locations by the example individual. However, for the purpose of recovering the combined anchor locations and road segments $\mathcal{A} \cup \mathcal{R}$, the level sets from density ranking are significantly more accurate compared to the level sets from kernel density estimation.

In Figure 7 we display the three summary curves presented in Section 5. Each panel contains 100 curves corresponding with each simulation replicate, but many of these curves overlap. The mass-volume curve and the Betti number curve (left and middle panels) are flat around the intervals $[0, 0.3]$, $[0.3, 0.5]$ and $[0.5, 0.6]$. These flat regions provide insight

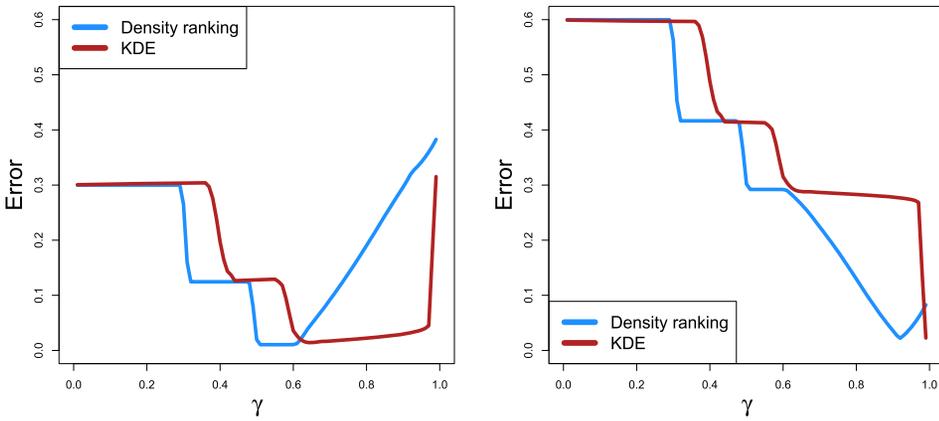


FIG. 6. Errors for detecting the anchor locations \mathcal{A} (left panel) and the combined anchor locations and road segments $\mathcal{A} \cup \mathcal{R}$ (right panel) in the simulation study. The x-axis shows the level γ . In the left panel, the y-axis shows the error $\text{PGPS}(A_\gamma \Delta \mathcal{A})$. In the right panel, the y-axis shows the error $\text{PGPS}(A_\gamma \Delta (\mathcal{A} \cup \mathcal{R}))$.

about the existence of anchor locations: $[0, 0.3]$ is for home, $[0.3, 0.5]$ is for work, $[0.5, 0.6]$ is for gym. The persistence curve (right panel) also indicates that there are three connected components with high persistences. Each connected component is associated with an anchor point.

The summary curves allow us to choose the density ranking level to recover the anchor points. For example, the flat region at $[0.5, 0.6]$ of the mass-volume curve and the Betti number curve correspond to the valley in the error curves in the left panel of Figure 6. As such, using level sets \hat{A}_γ with $\gamma \in [0.5, 0.6]$ to estimate \mathcal{A} yields the smallest estimation errors. Thus, the summary curves are very informative about the choice of ranking thresholds to employ in the identification of anchor points.

8. Analysis of GPS data. We illustrate the application of our methodology to the GPS data from the pilot study described in Section 2.

8.1. *Density ranking.* We apply density ranking based on the KDE in equation (2) with a smoothing bandwidth h of 200 meters. This choice implies that every observed GPS location will affect its neighborhood up to a distance of 200 meters. In the Supplementary Material (Chen and Dobra (2020)), Section 1, we compare several smoothing bandwidths: $h = 200$ seems to give an appropriate amount of smoothing for these data.

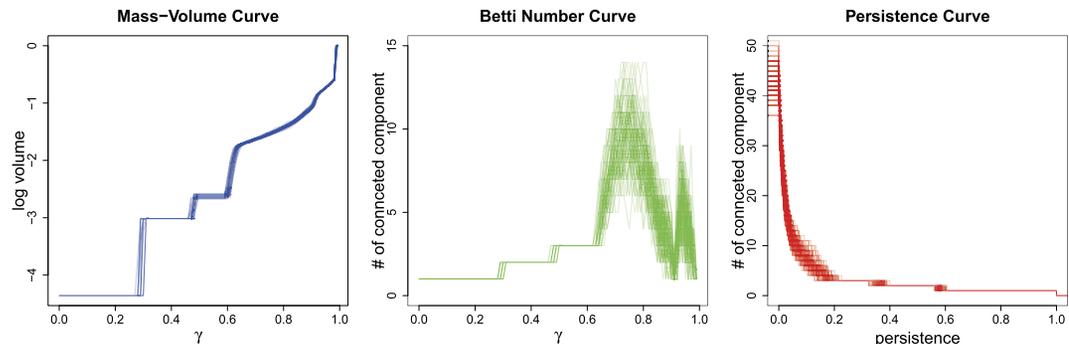


FIG. 7. Summary curves from the 100 simulated datasets. The left panel shows the mass-volume curve, the middle panel shows the Betti number curve and the right panel shows the persistence curve. The three flat regions in the left and middle panels correspond to the three anchor points \mathcal{A} .

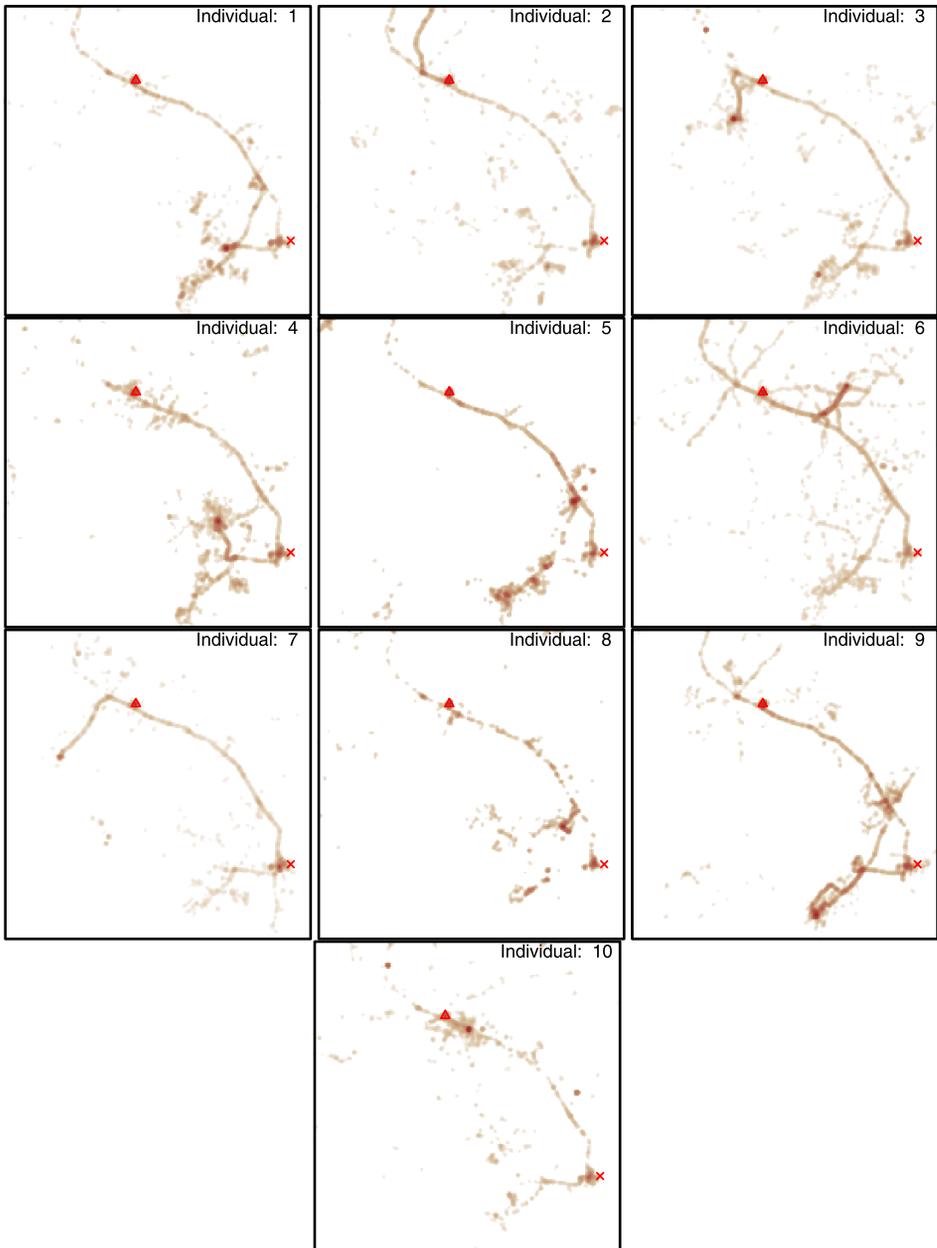


FIG. 8. Maps showing the density ranking of the locations recorded for each of participant in the GPS pilot study. The location of the workplace of the study participants is marked with a red triangle, while the location of the center of the township in the study area is marked with a red cross.

We consider GPS locations that belong to the zoom-in area shown in Figure 1 since this area contains most locations of the study participants. The density ranking of each study participant is shown in Figure 8. The pattern of density ranking varies from individual to individual. Study participants 2, 4 and 6 have more widespread GPS location distributions, while study participants 1 and 5 recorded GPS locations that seem to be more clustered. There are two key locations shared by all study participants: the workplace and the location of the center of the township. The density ranking of all study participants is high at the locations of the workplace and the township along the road that connects them.

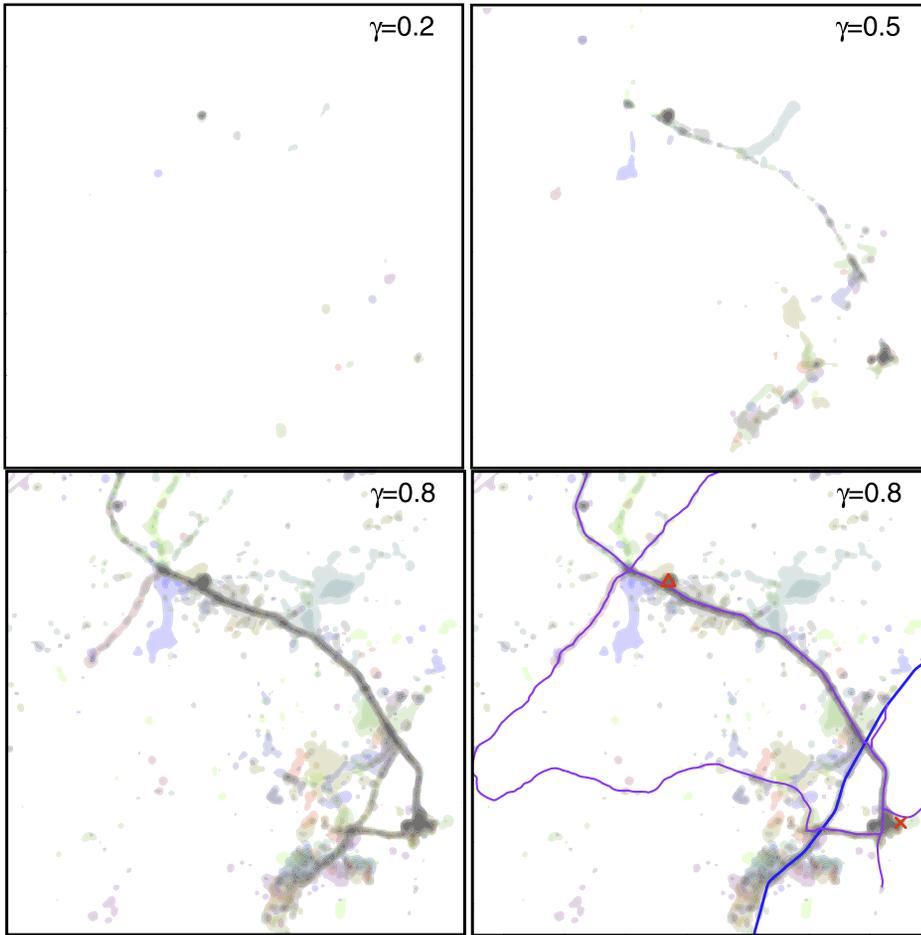


FIG. 9. The top activity spaces at levels $\gamma = 0.2, 0.5, 0.8$ from density ranking of the study participants. Colors indicate different individuals. The bottom right panel also shows the primary and secondary roads. There are two key locations shared by all study participants: The workplace (red triangle) and the location of the center of the township (red cross).

In Figure 9 we overlap the top activity spaces of the study participants at three levels: $\gamma = 0.2, 0.5$ and 0.8 . The workplace was included in all top 20% activity spaces (top left panel), while the township was included in all top 50% activity spaces (top right panel). Paths that follow several local roads are included in most of the top 80% activity spaces (bottom panels). Except for the workplace and the township, the rest of the top 20% activity spaces of the study participants are not overlapping: these regions are probably indicative of the locations of their homes.

8.2. Mass-volume curve. In Figure 10 we give the mass-volume curves of the study participants. We plot the function $\log V(\gamma)$ instead of $V(\gamma)$ since the size of activity space evolves rapidly when γ changes. The gray curve which corresponds to study participant 9 dominates the others in the range of $\gamma \in [0.1, 0.7]$, while the purple curve which corresponds to individual 6 takes over when $\gamma > 0.7$. This means that study participant 9 has the highest degree of mobility when we consider the activity space of top 10–70% activities. Individual 6 has the highest degree of mobility in terms of the activity space of top 70% or higher activities. The reason why these two curves dominate the others can be seen in Figure 8. The regions \hat{A}_γ for $\gamma \in [0.1, 0.7]$ correspond to where the density ranking is between 0.3–0.9

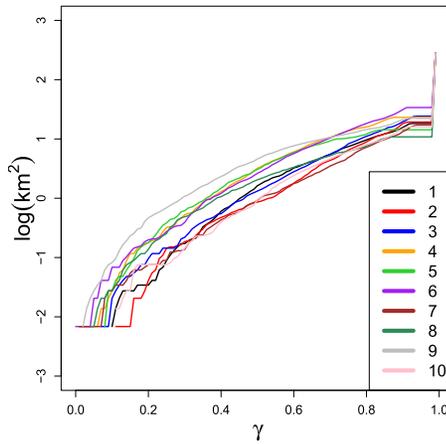


FIG. 10. *The mass-volume curves of the study participants measured on the log scale.*

$(1 - \gamma)$ which is the region with a darker color. The contours associated with study participant 9 have a wider region with darker color compared to others. When we consider regions with $\gamma > 0.7$, we are looking at regions with a lighter color. In this case we see that the density ranking of individual 6 spans a larger area compared to others. The mass-volume curves flatten out when the log size of the area is roughly below -2 . This is due to the resolution of the raster grid of cells used to compute the size of the level sets. The corresponding calculations cannot be performed if the size of the level sets falls below the resolution of the grid.

8.3. Betti number curve. Figure 11 shows the Betti number curves of the study participants. There are three curves that dominate the others for different ranges of γ . When $\gamma < 0.5$, the gray curve associated with study participant 9 dominates the others. When $0.5 < \gamma < 0.7$, the orange curve (study participant 4) dominates. When $0.7 < \gamma$, the purple curve (study participant 6) is the highest. This means that when we consider the activity spaces of top 50% activity (or an even higher level of activity), the activity space of study participant 9 has the largest number of connected components. This can actually be seen in Figure 8: the darker regions in density ranking of study participant 9 have more distinct connected components. The contours of density ranking of study participants 4 and 6 have many

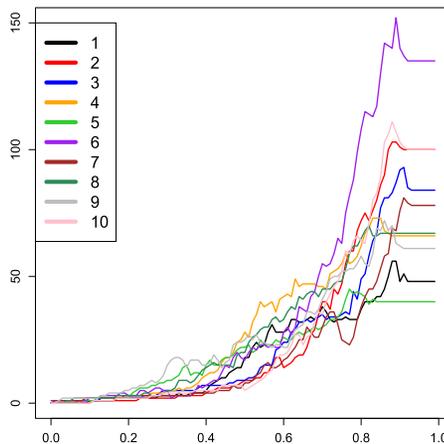


FIG. 11. *The Betti number curves of the study participants.*

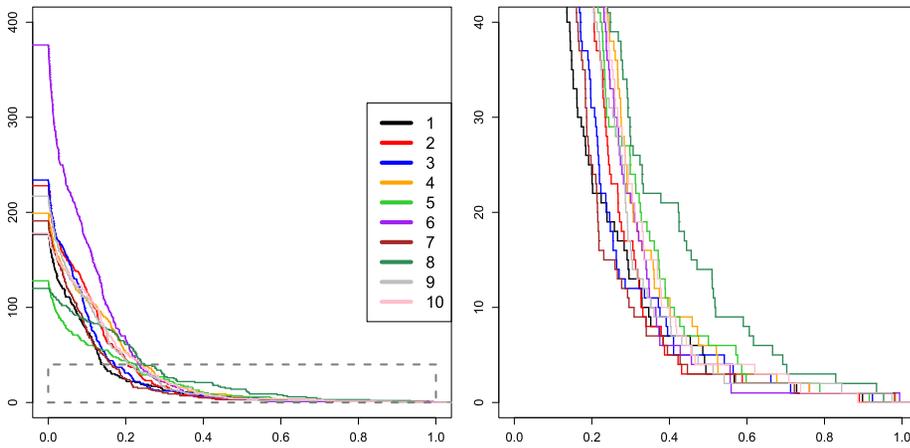


FIG. 12. *The persistence curves of the study participants. The left panel shows the full range of the persistence curve, and the right panel is the zoom-in version of the gray dashed box in the left panel.*

small bumps, resulting in a large number of connected components. The black and green curves associated with study participants 1 and 5 are smaller than 50. This is the result of their density ranking contours (Figure 8) being very concentrated. Unlike the density ranking of study participants 4, 6 and 9 that have many little bumps, the contours of study participants 1 and 5 do not have a spurious distribution which keeps their Betti number curves at lower values.

Based on Figure 11, we say that study participants 4, 6 and 9 have a higher degree of mobility, while study participants 1 and 5 have lower mobility in terms of the number of connected components of their activity spaces. Remark that the higher degree of mobility of study participant 4 becomes apparent based on the Betti number curve but is not evident based on the mass-volume curve.

8.4. Persistence curve. Figure 12 presents the persistence curves of the 10 study participants. The left panel displays the persistence curve in the full range, and the right panel is the zoom-in version with the range of y -axis restricted to the interval $[0, 40]$. In the left panel we see that the purple curve (study participant 6) dominates the others when the range of persistence is within $[0, 0.2]$. This range corresponds to many small bumps in its distribution of density ranking; see Figure 8. These small bumps create several connected components with short life spans: they all merge with other connected components quickly, so they have small persistence. These connected components with short life spans contribute to the larger values of the persistence curve. The fact that study participant 6 has many small and spurious bumps in their density ranking implies that this person repeatedly visits a larger number of locations. It is possible that this study participant has a job that involves driving on a daily basis.

In the right panel of Figure 12, the dark green curve that corresponds to study participant 8 stands out. This means that study participant 8 has more persistent connected components when we threshold on the persistence with a level above 0.4. This also implies that the activities of study participant 8 have several modes. From the density ranking distribution of study participant 8 (Figure 8), we see that this individual has several distinct connected components isolated from each other, confirming that this individual's activities have several modes. This is not evident from the mass-volume curve and the Betti number curve. Therefore, the persistence curve reveals key information about human mobility which complements the information provided by the other two types of curves we discussed.

9. Discussion. In this paper we described the key elements of human activity spaces (anchor locations, roads and areas around anchor locations), and proposed a mixture model for representing these elements. We discussed density ranking as an alternative to KDE, presented three types of summary curves and demonstrated their relevance for determining the geometry, size and structure of human activity spaces. We remark that these summary curves can also be calculated based on the KDE. However, using kernel density estimation instead of density ranking is not advisable since, as we proved in Section 6, the KDE's expectation diverges at anchor locations and along road segments. Density ranking has a powerful property that guarantees its convergence even when the underlying distribution contains lower dimensional structures (Chen (2019)). For this reason it is more appropriate to employ density ranking as opposed to KDE in the determination and measurement of human activity spaces from GPS data. As we illustrate in the Supplementary Material (Chen and Dobra (2020)), Section 2, density ranking is preferable to KDE when creating maps of GPS locations: finding an appropriate transformation of the KDE that will help create an informative map of an activity pattern could be challenging. Transforming the KDE into density ranking preserves the relevant information in the location data and offers a simple, generally applicable solution to this problem.

The collection of high resolution movement data of individuals over long periods of time is possible thanks to today's technological advances. Smartphones are an especially versatile device that an ever-growing proportion of people from most countries carry around every day. At the present time GPS datasets collected from smartphones are recorded as part of federally funded studies from many research fields. This collection effort will, without doubt, continue to expand in the coming years and will provide detailed information about where people spend their time. The methods we presented in this paper could constitute a key component of these studies that will help translate raw GPS locations into meaningful, easily interpretable information about individuals' daily selective mobility. We demonstrated that density ranking and summary curves have substantial advantages over existent methods for activity space determination since they are not constrained to a fixed geometrical shape, allow the determination of anchor locations and roads used for travel, are less influenced by outlier locations and are not dependent on the availability of quality road network data.

Human activity spaces are fundamental for health research (Perchoux et al. (2013)), and can be interpreted as indicators of social activity, self-confidence and knowledge about the physical environment. They capture the dynamics of the geographic context (Kwan (2012)) which is critical in assessing individuals' exposure to social and environmental risk factors over multiple neighborhoods that are visited during activities of daily living. In particular, they are one of the foundation constructs of contextual expology (Chaix et al. (2012), Kwan (2009)). This is a subdiscipline that focuses on modeling the individuals' spatiotemporal patterns of exposure and on the derivation of related multiplace environmental exposure variables. The premise is that even individuals from the same residential community could spend different amounts of time away from their home and travel to locations with different characteristics. This leads to various levels of exposure to spatially-varying risk factors. Contextual expology creates customized exposure measures based on the shape, spatial spread and configuration of the activity space of each person by taking into account their spatial polygamy (Matthews (2008, 2011)), that is, the amount of time spent at, around or traveling between their anchor locations. Density ranking and summary curves could be used in developing much needed exposure measures to contextual or environmental influences that take into account the spatiotemporal patterns of human mobility (Kwan (2013)).

An open research question relates to linking sociodemographic characteristics of individuals with their activity spaces and studying the interactions that might exist between the characteristics of places and the characteristics of individuals that visit these places (Schönfelder

and Axhausen (2003)). Research on activity spaces could lead to more effective individual-tailored interventions that take into consideration multiple geographic contexts. Such interventions could provide customized information to individuals about sources of healthy food, outdoor places to walk or to exercise or local social events based on their own spatial mobility patterns.

Acknowledgments. We thank two referees for the very helpful comments on this paper. Yen-Chi Chen received partial support from the National Science Foundation Grant DMS-1810960 and National Institutes of Health Grant U01-AG016976. Adrian Dobra received partial support from the National Science Foundation Grant DMS/MPS-1737746. The funders had no role in the study design, data collection and analysis, decision to publish or preparation of the manuscript.

SUPPLEMENTARY MATERIAL

Additional proofs, explanations and figures (DOI: [10.1214/19-AOAS1311SUPP](https://doi.org/10.1214/19-AOAS1311SUPP); .pdf). In this online Supplementary Material (Chen and Dobra (2020)), we give the proofs of the theoretical results presented in the paper. We also provide additional explanations and numerical results.

REFERENCES

- APOSTOLOPOULOS, Y. and SONMEZ, S. (2007). *Population Mobility and Infectious Disease*. Springer, New York.
- BASTA, L. A., RICHMOND, T. S. and WIEBE, D. J. (2010). Neighborhoods, daily activities, and measuring health risks experienced in urban environments. *Soc. Sci. Med.* **71** 1943–1950.
- BERRIGAN, D., HIPPI, J. A., HURVITZ, P. M., JAMES, P., JANKOWSKA, M. M., KERR, J., LADEN, F., LEONARD, T., MCKINNON, R. A. et al. (2015). Geospatial and contextual approaches to energy balance and health. *Ann. of GIS* **21** 157–168.
- BERRY, E., CHEN, Y.-C., CISEWSKI-KEHE, J. and FASY, B. T. (2018). Functional summaries of persistence diagrams. Preprint. Available at [arXiv:1804.01618](https://arxiv.org/abs/1804.01618).
- BISCIO, C. A. N. and MØLLER, J. (2019). The accumulated persistence function, a new useful functional summary statistic for topological data analysis, with a view to brain artery trees and spatial point process applications. *J. Comput. Graph. Statist.* **28** 671–681. MR4007749 <https://doi.org/10.1080/10618600.2019.1573686>
- BULIUNG, R. N. (2001). Spatiotemporal patterns of employment and non-work activities in Portland, Oregon. In *ESRI International User Conference*, San Diego, CA.
- BULIUNG, R. N. and KANAROGLOU, P. S. (2006). Urban form and household activity-travel behavior. *Growth Change* **37** 172–199.
- CADRE, B., PELLETIER, B. and PUDLO, P. (2013). Estimation of density level sets with a given probability content. *J. Nonparametr. Stat.* **25** 261–272. MR3039981 <https://doi.org/10.1080/10485252.2012.750319>
- CARLSSON, G. (2009). Topology and data. *Bull. Amer. Math. Soc. (N.S.)* **46** 255–308. MR2476414 <https://doi.org/10.1090/S0273-0979-09-01249-X>
- CHAIX, B., KESTENS, Y., PERCHOUX, C., KARUSISI, N., MERLO, J. and LABADI, K. (2012). An interactive mapping tool to assess individual mobility patterns in neighborhood studies. *Am. J. Prev. Med.* **43** 440–450.
- CHAZAL, F. and MICHEL, B. (2017). An introduction to topological data analysis: Fundamental and practical aspects for data scientists. Preprint. Available at [arXiv:1710.04019](https://arxiv.org/abs/1710.04019).
- CHEN, Y.-C. (2019). Generalized cluster trees and singular measures. *Ann. Statist.* **47** 2174–2203. MR3953448 <https://doi.org/10.1214/18-AOS1744>
- CHEN, Y.-C. and DOBRA, A. (2020). Supplement to “Measuring human activity spaces from GPS data with density ranking and summary curves.” <https://doi.org/10.1214/19-AOAS1311SUPP>.
- CHEN, C., MA, J., SUSILO, Y., LIU, Y. and WANG, M. (2016). The promises of big data and small data for travel behavior (aka human mobility) analysis. *Transp. Res., Part C, Emerg. Technol.* **68** 285–299.
- CHRISTIAN, W. J. (2012). Using geospatial technologies to explore activity-based retail food environments. *Spatial and Spatio-Temporal Epidemiology* **3** 287–295.
- CLÉMENÇON, S. and JAKUBOWICZ, J. (2013). Scoring anomalies: A m-estimation formulation. In *Artificial Intelligence and Statistics* 659–667.

- CLÉMENÇON, S. and THOMAS, A. (2018). Mass volume curves and anomaly ranking. *Electron. J. Stat.* **12** 2806–2872. MR3855357 <https://doi.org/10.1214/18-EJS1474>
- CUMMINS, S., CURTIS, S., DIEZ-ROUX, A. V. and MACINTYRE, S. (2007). Understanding and representing ‘place’ in health research: A relational approach. *Soc. Sci. Med.* **65** 1825–1838.
- DOBRA, A., WILLIAMS, N. E. and EAGLE, N. (2015). Spatiotemporal detection of unusual human population behavior using mobile phone data. *PLoS ONE* **10** 1–20.
- DOBRA, A., BÄRNIGHAUSEN, T., VANDORMAEL, A. and TANSER, F. (2017). Space-time migration patterns and risk of HIV acquisition in rural South Africa. *AIDS* **31** 137–145.
- DOSS, C. R. and WENG, G. (2018). Bandwidth selection for kernel density estimators of multivariate level sets and highest density regions. *Electron. J. Stat.* **12** 4313–4376. MR3892342 <https://doi.org/10.1214/18-ejs1501>
- EDELSBRUNNER, H. and HARER, J. (2008). Persistent homology—a survey. In *Surveys on Discrete and Computational Geometry. Contemp. Math.* **453** 257–282. Amer. Math. Soc., Providence, RI. MR2405684 <https://doi.org/10.1090/conm/453/08802>
- EDELSBRUNNER, H. and HARER, J. L. (2010). *Computational Topology: An Introduction*. Amer. Math. Soc., Providence, RI. MR2572029
- ENTWISLE, B. (2007). Putting people into place. *Demography* **44** 687–703.
- FAN, Y. and KHATTAK, A. (2008). Urban form, individual spatial footprints, and travel: Examination of space-use behavior. *Transp. Res. Rec.* **2082** 98–106.
- GHRIST, R. (2008). Barcodes: The persistent topology of data. *Bull. Amer. Math. Soc. (N.S.)* **45** 61–75. MR2358377 <https://doi.org/10.1090/S0273-0979-07-01191-3>
- GHRIST, R. (2014). *Elementary Applied Topology*. CreateSpace Independent Publishing Platform.
- GOLLEDGE, R. G. (1999). Human wayfinding and cognitive maps. In *Wayfinding Behavior* (R. G. Golledge, ed.) 5–45. The Johns Hopkins Univ. Press, Baltimore, MD.
- GOLLEDGE, R. G. and STIMSON, R. J. (1997). *Spatial Behavior*. The Guildford Press, New York.
- HÄGERSTRAND, T. (1963). Geographic measurements of migration. In *Human Displacements: Measurement Methodological Aspects* (J. Sutter, ed.), Monaco.
- HÄGERSTRAND, T. (1970). What about people in regional science? *Pap. Reg. Sci.* **24** 7–21.
- HARDING, C., PATTERSON, Z. and MIRANDA-MORENO, L. F. (2013). Activity space geometry and its effect on mode choice. In *Transportation Research Board 92nd Annual Meeting*, Washington DC.
- HARTIGAN, J. A. (1975). *Clustering Algorithms. Wiley Series in Probability and Mathematical Statistics*. Wiley, New York. MR0405726
- HIRSCH, J. A., WINTERS, M., CLARKE, P. and MCKAY, H. (2014). Generating GPS activity spaces that shed light upon the mobility habits of older adults: A descriptive analysis. *Int. J. Health Geogr.* **13** 51.
- HURVITZ, P. M., MOUDON, A. V., KANG, B., SAELENS, B. E. and DUNCAN, G. E. (2014). Emerging technologies for assessing physical activity behaviors in space and time. *Front. Public Health* **2**.
- KACZYNSKI, T., MISCHAIKOW, K. and MROZEK, M. (2004). *Computational Homology. Applied Mathematical Sciences* **157**. Springer, New York. MR2028588 <https://doi.org/10.1007/b97315>
- KESTENS, Y., LEBEL, A., DANIEL, M., THÉRIAULT, M. and PAMPALON, R. (2010). Using experienced activity spaces to measure foodscape exposure. *Health Place* **16** 1094–1103.
- KIM, S. and ULFARSSON, G. F. (2015). Activity space of older and working-age adults in the Puget Sound region. In *Transportation Research Board 94th Annual Meeting*, Washington DC.
- KWAN, M. P. (2000). Interactive geovisualization of activity-travel patterns using three-dimensional geographical information systems: A methodological exploration with a large data set. *Transp. Res.* **8C** 185–203.
- KWAN, M.-P. (2009). From place-based to people-based exposure measures. *Soc. Sci. Med.* **69** 1311–1313.
- KWAN, M.-P. (2012). The uncertain geographic context problem. *Ann. Assoc. Am. Geogr.* **102** 958–968.
- KWAN, M.-P. (2013). Beyond space (as we knew it): Toward temporally integrated geographies of segregation, health, and accessibility. *Ann. Assoc. Am. Geogr.* **103** 1078–1086.
- LEE, J. H., DAVIS, A. W., YOON, S. Y. and GOULIAS, K. G. (2016). Activity space estimation with longitudinal observations of social media data. *Transp.* **43** 955–977.
- LUM, P. Y., SINGH, G., LEHMAN, A., ISHKANOV, T., VEJDEMO-JOHANSSON, M., ALAGAPPAN, M., CARLSSON, J. and CARLSSON, G. (2013). Extracting insights from the shape of complex data using topology. *Sci. Rep.* **3** 1236.
- MANAUGH, K. and EL-GENEIDY, A. M. (2012). What makes travel ‘local’: Defining and understanding local travel behavior. *J. Transp. Land Use* **5** 12–27.
- MASON, D. M. and POLONIK, W. (2009). Asymptotic normality of plug-in level set estimates. *Ann. Appl. Probab.* **19** 1108–1142. MR2537201 <https://doi.org/10.1214/08-AAP569>
- MATTHEWS, S. A. (2008). The salience of neighborhood. *Am. J. Prev. Med.* **34** 257–259.
- MATTHEWS, S. A. (2011). Spatial polygamy and the heterogeneity of place: Studying people and place via ego-centric methods. In *Communities, Neighborhoods, and Health: Expanding the Boundaries of Place* (L. M. Burton, S. A. Matthews, M. C. Leung, S. P. Kemp and D. T. Takeuchi, eds.) 35–55. Springer, New York, NY.

- MATTHEWS, S. A. and YANG, T.-C. (2013). Spatial polygamy and contextual exposures (SPACES): Promoting activity space approaches in research on place and health. *Am. Behav. Sci.* **57** 1057–1081.
- MATTILA, P. (1995). *Geometry of Sets and Measures in Euclidean Spaces: Fractals and Rectifiability*. Cambridge Studies in Advanced Mathematics **44**. Cambridge Univ. Press, Cambridge. MR1333890 <https://doi.org/10.1017/CBO9780511623813>
- NEWSOME, T. H., WALCOTT, W. A. and SMITH, P. D. (1998). Urban activity spaces: Illustrations and application of a conceptual model for integrating the time and space dimensions. *Transp.* **25** 357–377.
- NUÑEZ GARCIA, J., KUTALIK, Z., CHO, K.-H. and WOLKENHAUER, O. (2003). Level sets and minimum volume sets of probability density functions. *Internat. J. Approx. Reason.* **34** 25–47. MR2017778 [https://doi.org/10.1016/S0888-613X\(03\)00052-5](https://doi.org/10.1016/S0888-613X(03)00052-5)
- PERCHOUX, C., CHAIX, B., CUMMINS, S. and KESTENS, Y. (2013). Conceptualization and measurement of environmental exposure in epidemiology: Accounting for activity space related to daily mobility. *Health Place* **21** 86–93.
- POLONIK, W. (1997). Minimum volume sets and generalized quantile processes. *Stochastic Process. Appl.* **69** 1–24. MR1464172 [https://doi.org/10.1016/S0304-4149\(97\)00028-8](https://doi.org/10.1016/S0304-4149(97)00028-8)
- PREISS, D. (1987). Geometry of measures in \mathbb{R}^n : Distribution, rectifiability, and densities. *Ann. of Math. (2)* **125** 537–643. MR0890162 <https://doi.org/10.2307/1971410>
- QIAO, W. (2017). Asymptotics and optimal bandwidth selection for nonparametric estimation of density level sets. Preprint. Available at [arXiv:1707.09697](https://arxiv.org/abs/1707.09697).
- RAI, R. K., BALMER, M., RIESER, M., VAZE, V. S., SCHÖNFELDER, S. and AXHAUSEN, K. W. (2007). Capturing human activity spaces: New geometries. *Transp. Res. Rec.* **2021** 70–80.
- RICHARDSON, D. B., VOLKOW, N. D., KWAN, M.-P., KAPLAN, R. M., GOODCHILD, M. F. and CROYLE, R. T. (2013). Spatial turn in health research. *Science* **339** 1390–1392.
- RIGOLLET, P. and VERT, R. (2009). Optimal rates for plug-in estimators of density level sets. *Bernoulli* **15** 1154–1178. MR2597587 <https://doi.org/10.3150/09-BEJ184>
- RINALDO, A. and WASSERMAN, L. (2010). Generalized density clustering. *Ann. Statist.* **38** 2678–2722. MR2722453 <https://doi.org/10.1214/10-AOS797>
- SCHÖNFELDER, S. and AXHAUSEN, K. W. (2003). Activity spaces: Measures of social exclusion? *Transp. Policy* **10** 273–286.
- SCHÖNFELDER, S. and AXHAUSEN, K. W. (2004). On the variability of human activity spaces. In *The Real and Virtual Worlds of Spatial Planning* (M. Koll-Schretzenmayr, M. Keiner and G. Nussbaumer, eds.) 237–262. Springer, Berlin, Heidelberg.
- SCOTT, C. D. and NOWAK, R. D. (2006). Learning minimum volume sets. *J. Mach. Learn. Res.* **7** 665–704. MR2274383
- SHERMAN, J. E., SPENCER, J., PREISSER, J. S., GESLER, W. M. and ARCURY, T. A. (2005). A suite of methods for representing activity space in a healthcare accessibility study. *Int. J. Health Geogr.* **4** 24–24.
- SILVERMAN, B. W. (1986). *Density Estimation for Statistics and Data Analysis*. Monographs on Statistics and Applied Probability. CRC Press, London. MR0848134 <https://doi.org/10.1007/978-1-4899-3324-9>
- WASSERMAN, L. (2016). Topological data analysis. Available at [arXiv:1609.08227](https://arxiv.org/abs/1609.08227).
- WASSERMAN, L. (2018). Topological data analysis. *Annu. Rev. Stat. Appl.* **5** 501–535. MR3774757 <https://doi.org/10.1146/annurev-statistics-031017-100045>
- WIEHE, S. E., CARROLL, A. E., LIU, G. C., HABERKORN, K. L., HOCH, S. C., WILSON, J. S. and FORTENBERRY, J. D. (2008). Using gps-enabled cell phones to track the travel patterns of adolescents. *Int. J. Health Geogr.* **7** 22–22.
- WIEHE, S., KWAN, M.-P., WILSON, J. and FORTENBERRY, J. (2013). Adolescent health-risk behavior and community disorder. *PLoS ONE* **8** e77667.
- WILLIAMS, N. E., THOMAS, T. A., DUNBAR, M., EAGLE, N. and DOBRA, A. (2015). Measures of human mobility using mobile phone records enhanced with gis data. *PLoS ONE* **10** 1–16.
- WONG, D. W. S. and SHAW, S.-L. (2011). Measuring segregation: An activity space approach. *J. Geogr. Syst.* **13** 127–145.
- WORTON, B. J. (1987). A review of models of home range for animal movement. *Ecol. Model.* **38** 277–298.
- ZENK, S. N., SCHULZ, A. J., MATTHEWS, S. A., ODOMS-YOUNG, A., WILBUR, J., WEGRZYN, L., GIBBS, K., BRAUNSCHWEIG, C. and STOKES, C. (2011). Activity space environment and dietary and physical activity behaviors: A pilot study. *Health Place* **17** 1150–1161.
- ZENK, S. N., SCHULZ, A. J., ODOMS-YOUNG, A., WILBUR, J., MATTHEWS, S. A., GAMBOA, C., WEGRZYN, L. R., HOBSON, S. and STOKES, C. (2012). Feasibility of using global positioning systems (GPS) with diverse urban adults: Before and after data on perceived acceptability, barriers, and ease of use. *J. Phys. Act. Health* **9** 924–934.