



# Understanding the dynamics of urban areas of interest through volunteered geographic information

Meixu Chen<sup>1</sup> · Dani Arribas-Bel<sup>1</sup> · Alex Singleton<sup>1</sup>

Received: 10 October 2017 / Accepted: 14 November 2018 / Published online: 5 December 2018  
© The Author(s) 2018

## Abstract

Obtaining insights about the dynamics of urban structure is crucial to the framing of the context within the smart city. This paper focuses on urban areas of interest (UAOI), a concept that provides functional definitions of a city's spatial structure. Traditional sources of social data can rarely capture these aspects at scale while spatial information on the city alone does not capture how the population values different parts of the city and in different ways. Hence, we leverage volunteered geographic information (VGI) to overcome some of the limits of traditional sources in providing urban structural and functional insights. We use a special type of VGI—metadata from geotagged Flickr images—to identify UAIOs and exploit their temporal and spatial attributes. To do this, we propose a methodological strategy that combines hierarchical density-based spatial clustering for applications with noise and the 'α-shape' algorithm to quantify the dynamics of UAIOs in Inner London for a period 2013–2015 and develop an innovative visualisation of UAIO profiles from which UAIO dynamics can be explored. Our results expand and improve upon the previous literature on this topic and provide a useful reference for urban practitioners who might wish to include more timely information when making decisions.

**Keywords** Urban dynamics · Urban areas of interest · Quantitative analysis · Volunteered geographic information · Social media data

**JEL Classification** R0

---

✉ Meixu Chen  
meixu@liverpool.ac.uk

<sup>1</sup> Geographic Data Science Lab, Department of Geography and Planning, University of Liverpool, Roxby Building, 74 Bedford St S, Liverpool L69 7ZT, UK

## 1 Introduction

The rapid growth of urban populations across the globe is resulting in new kinds of technical, physical, material and social challenges and constraints (Chourabi et al. 2012). With the aim of tackling such issues, how to make a city ‘smart’ has become a significant strategy in many developed and developing regions of the world. Although there is no standard definition of the smart city, common characteristics can be summarised as an integrated system connecting digital technologies, critical infrastructures and citizens, to plan, govern and manage a city in order to improve its sustainability, optimise processes and maximise the provision of collective public and private services (Harrison et al. 2010; Washburn et al. 2010; Batty 2017). Integral to the operationalisation of a smart city, it is often of relevance to obtain timely insights into the dynamics of urban population at a temporal granularity finer than that of traditional surveys, which can be enhanced by or provided through digital technologies.

It is within this context that the present paper engages with the concept of urban areas of interest (UAOI), which refers to parts of the urban built environment that can be delineated in their extent through the clustering of human activity. Such areas may contain business zones, tourist attractions, iconic landmarks, recreational zones or other attractors (Hu et al. 2015). The notion of a UAOI is, therefore, a combination of morphological features including buildings and streets, and ‘points of interest’, as defined by the relevance the population concedes specific parts of cities. As such, a UAOI can be viewed as a perceptual space, which is captured by the social morphology of the city, albeit rooted in physical space (Crooks et al. 2016). Accordingly, a UAOI should emerge from the activities of a large collective of different people to avoid very individual conceptions. Furthermore, such definitions are complex, as unlike well-defined geographic divisions or administrative districts, the delineation of a UAOI may vary between people in different contexts, ages and cultures.

Identifying and understanding UAOIs has applications in multiple fields. For spatial planning, they may assist in identifying areas with greater public priority in the context of limited resource availability (Gandy 2006). For retailing, they can help identify areas where people cluster, and how these have evolved over time, which might aid in store location or for targeting advertisements more effectively. For transport planning, they may help prioritise traffic flows or the provision of public transport; for statistical agencies, they may provide useful reference distributions in comparison with official geographical divisions.

The challenge of defining UAOIs over time resides in the need for granular spatiotemporal data recorded within cities. Although traditional data sources used in urban studies, such as remotely sensed data, have a lengthy history of application and can be used to characterise urban morphology, they do not capture human dynamics beyond expansion or contraction of the built form. Alternatively, survey or census data might be utilised to inform the discovery of UAOIs, but these are usually costly to administer and may be of limited temporal granularity (Shi et al. 2014; Tasse and Hong 2014). A third alternative has emerged in the last few years. Several new forms of digital data

derived from urban activity through passive or active forms of data collection capture urban form and/or social functional geography (Arribas-Bel 2014; Crooks et al. 2016). Such data are referred to as volunteered geographic information (VGI; Goodchild 2007), which includes the use of digital devices by communities or individuals to create, accumulate, upload and communicate geographic information, typically through contemporary web technology. Commonly designated as VGI is a variety of content from social media networks, which often support geolocation of assets and include networks such as Twitter, Facebook, Flickr and Instagram. Data derived from these networks have been used in a variety of contexts to explore spatial, temporal and even semantic information about human activities (Jiang et al. 2015; Lansley and Longley 2016; Lloyd and Cheshire 2017; Gao et al. 2017).

In this paper, we examine the potential of data derived from the online photograph management and sharing website Flickr to extract and understand urban areas of interest. Although there are inherent biases associated with geotagged Flickr data, a number of studies have utilised these data effectively to explore various issues within urban contexts (Hollenstein and Purves 2010; Lee et al. 2014; Hu et al. 2015; Gao et al. 2017). Flickr offers an attractive proposition as a data source for a number of reasons. The scale of the Flickr network is extensive and, as of 2016, Flickr had 122 million users with more than 10 billion images shared, demonstrating a large degree of penetration (Smith 2016). Secondly, the metadata of each Flickr photograph is available through its public application programming interface, which can be retrieved back to 2004, making it possible to consider the temporal dimension of imagery. These features are in contrast to other sources of VGI from social networks, which have rather limited data retrieval limits (e.g. Foursquare only allows 1 month; Foursquare for Developer 2017). Finally, studies have suggested that Flickr photographs, in most cases, are taken in the urban built environment, and as such, enhance their suitability as a source to identify UAOIs (Crandall et al. 2009; Hollenstein and Purves 2010).

Our goal is, therefore, to present a new method of extracting UAOIs and to provide new insights about their fine-grained spatiotemporal evolution and characteristics. We used geotagged Flickr data from three recent years (2013–2015) and have focused particularly on the seasonal variability of the UAOIs. A recent hierarchical algorithm was used to extract clusters, reducing many of the drawbacks of traditional, previously used and density-based methods. An ‘ $\alpha$ -shape’ algorithm was then utilised to construct boundaries identifying the UAOI extents. Once built, we conduct further analysis on the spatial and temporal patterns associated with the identified UAOIs and propose an approach to build a spatiotemporal profile for each UAOI.

The structure of this paper is organised as follows. The next section discusses related work about points of interest and areas of interest, as well as techniques for analysis of geotagged photograph data. Section 3 describes the data collection, data bias and pre-processing stages. Section 4, the core of the paper, proposes a methodological framework to extract and understand UAOIs, including an approach to validate the number of Flickr users in the extracted UAOIs. This is followed by a discussion of the spatiotemporal dynamics of the identified UAOIs. Finally, Sect. 6 concludes the paper and suggests future extensions to this research.

## 2 Literature review

There is a growing body of research that uses geotagged photographs, which have examined both the attributes contained within the metadata and the image itself. Most studies have focused on exploring landmark detection involving travel route recommendations, which generally integrate some aspect of movement/trajec-tory analysis (Zheng et al. 2012; Sun et al. 2015). Alternative approaches examine geotagged photographs to address the question of where and when events take place (Rattenbury et al. 2007; Kisilevich et al. 2010b; Papadopoulos et al. 2011). There is also further work that combines both image analyses with exploration of the meta-data, and applies it across a range of topics including detecting cultural differences (Yanai et al. 2009), land cover classification and validation (Antoniou et al. 2016) and definition of significant places on the basis of people's interaction with their surroundings (Li and Goodchild 2012).

The most prevalent use of how geotagged social media are used to extract points of interest are based upon exploiting the locational aspect of semantics (e.g. crowd-sourced tagging). For example, Crandall et al. (2009) presented techniques that can automatically identify popular places through representative images and textual labels from Flickr. Lee et al. (2014) proposed a framework to extract points of interest and their agglomerations from geotagged photographs. Andrienko and Andrienko (2013) extended such work through the additional consideration of the time of geotagged social media through a number of space–time visual analytic approaches. Other related work has extended the use of similar analysis techniques to include the exploration of attractive regions. For example, Kisilevich et al. (2010a) proposed a systematic framework for the exploration of points of interest obtained from Flickr and Panoramio, utilising a convex hull to create boundaries of concentrated areas for visualisation. Hollenstein and Purves (2010) also linked the derivation of data-driven density surfaces to the extraction of urban boundaries; this was extended by Li et al. (2013) who constructed spatial boundaries using kernel density estimation, which was utilised to approximate the number of place occurrences per unit area. In some sense, a generality between all of these was an aim of creating clusters of geotagged data. However, with limited exception, this line of inquiry has rarely focused on spatiotemporal changes, and acutely so over a multi-year period. Furthermore, although the popular non-parametric density estimation technique—kernel density estimation (KDE)—has examples of use to construct and visualise attractive aggregations of points of interest, this approach is not designed to delineate specific boundary lines of clusters, a valuable and necessary feature when identifying areas of interest.

One alternative to KDE is density-based clustering algorithms. This family of techniques has more recently been applied to identify points of interest or attractive areas (Kisilevich et al. 2010a, b; Lee et al. 2014; Andrienko and Andrienko 2013; Gao et al. 2017). The most widely used approach in this category is DBSCAN (density-based spatial clustering for applications with noise, Ester et al. 1996), which involves two parameters: the search radius (epsilon) and the minimum number of points (MinPts). Once both are specified, the algorithm identifies clusters of at least MinPts observations using epsilon as the maximum distance for the neighbour

search. However, both parameters need to be finely tuned, typically requiring manual experimentation in both cases before an appropriate value can be selected. In addition, DBSCAN only uses a global (single) density threshold to extract a flat partition, which fails to distinguish clusters of different densities. The OPTICS (ordering points to identify the clustering structure, Ankerst et al. 1999) algorithm presents an improvement to DBSCAN as it only requires the MinPts parameter to be specified while also producing a hierarchical result. However, this approach still relies on a global density threshold, which is unable to find the most significant clusters based on different density levels (Campello et al. 2013).

A recent application of DBSCAN with particular relevance to this paper was proposed by Hu et al. (2015), who presented a methodological approach that extracts UAOIs for six cities based on 10 years' worth of geotagged Flickr photographs. Building on this work, our research provides new insights into spatiotemporal changes in UAOIs by proposing a number of extensions. First, we focus on finer temporal scales, which allows us to consider seasonal variability. We demonstrate that this degree of resolution matters because it can capture seasonal UAOIs that emerge and disappear rapidly. Secondly, in terms of UAOI discovery, we introduce a more advanced method (Campello et al. 2013) than DBSCAN called hierarchical density-based spatial clustering for applications with noise (HDBSCAN) for extracting UAOIs. As discussed later in the methodological framework section, this approach overcomes some of the main drawbacks of other density-based clustering methods. Third, we propose the creation of spatiotemporal profiles based on small-scale geographic areas and use these to quantify the characteristics of spatiotemporal change in the UAOIs.

## 3 Data

### 3.1 Data description

Greater London was used as the study area because the regional boundaries contain a very large volume of geotagged photographs. Flickr data can be retrieved and downloaded using a public application programming interface (API, <https://www.flickr.com/services/api/>) through the Python interface (Stüvel 2016). Among the 10 billion images shared on Flickr, 3.33% contain geographic information (Smith 2016; Catt 2009). We used a bounding box to collect all geotagged data uploaded for Greater London. Dates between 1 January 2013 and 31 December 2015 were selected, as these 3 years have the highest number of Flickr photographs since Flickr was launched (Michel 2017). As this study focuses on the geo-temporal exploration of UAOIs, only locational and temporal metadata were retrieved and used in this study.<sup>1</sup> Our data set contained a total of 1,575,200 entries contributed by 34,615

<sup>1</sup> The use of only location and timestamp does not mean we do not recognise the value of other forms of metadata such as tags and text, or even content from each photo itself. These are fundamentally different sources of information that do not directly relate to the identification of UAOIs, and is hence beyond the scope of this paper. Their use is illustrated, for example, in Yanai et al. (2009), Crandall et al. (2009), Li et al. (2013) and Gao et al. (2017).

**Table 1** A sample of georeferenced Flickr metadata in London

User ID	Latitude	Longitude	Date
97938415@N00	51.507577	-0.099349	2015-01-19 19:40:42
30826685@N06	51.500504	-0.127419	2015-03-07 15:36:50
58216017@N02	51.499434	-0.163905	2013-11-29 16:00:27
26096388@N00	51.51353	-0.113193	2013-03-13 12:30:09
81434693@N00	51.500461	-0.138487	2013-03-09 12:23:55
89972047@N05	51.516541	-0.097525	2013-01-21 08:06:11
26602223@N00	51.530199	-0.125688	2014-03-30 15:27:43
10287726@N02	51.55811	-0.282823	2013-01-26 21:37:44

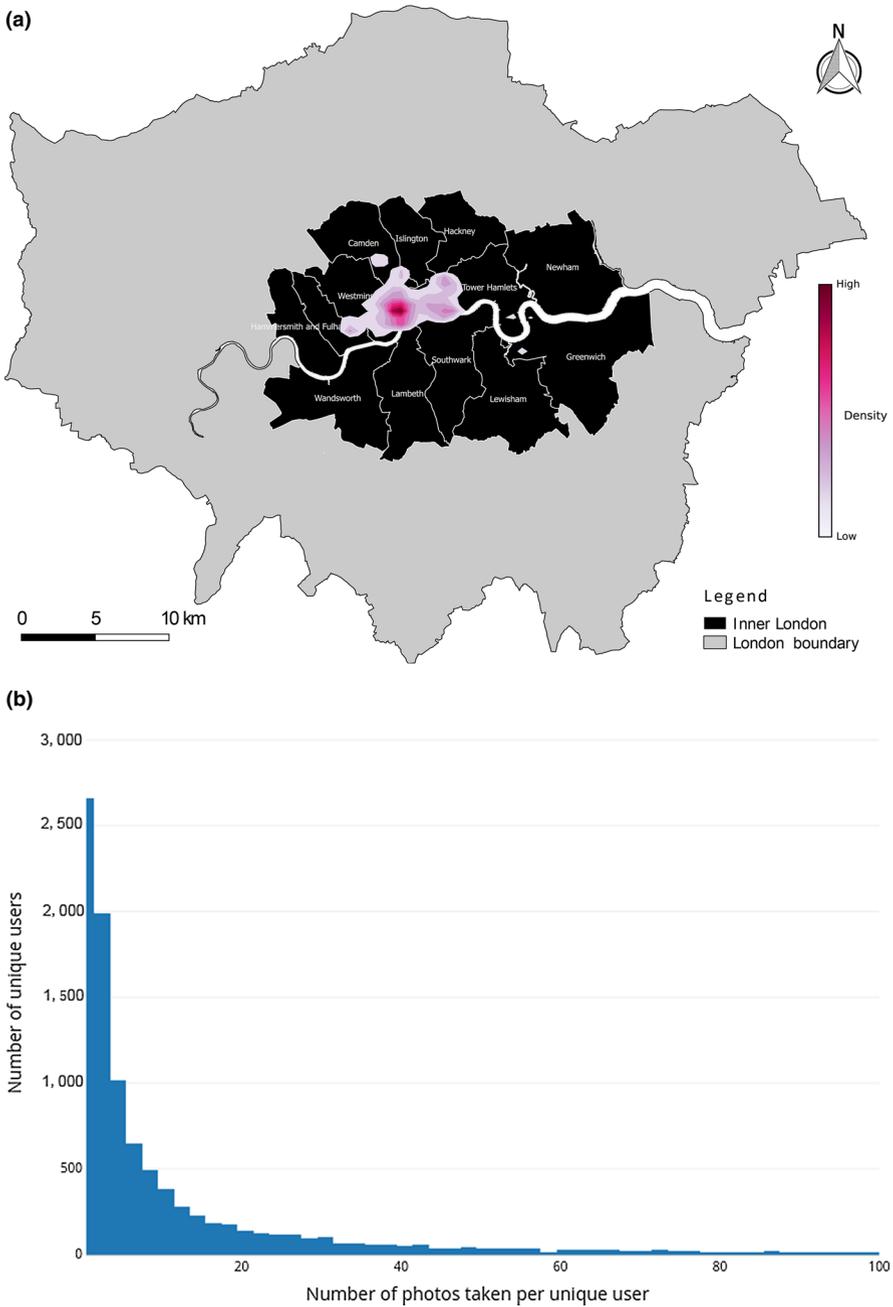
unique users, with the following attributes: user ID, geographic coordinates and timestamp (i.e. the time when the photograph was taken). Table 1 displays an extraction from the data.

As with other geotagged social media data, the Flickr sample had a few issues related to data quality. Data quality has been defined as data that are fit for use by data consumers (Wang and Strong 1996). In our case, “fit for use” involves allowing us to identify the spatial dynamics of the urban population. Data quality cannot be assured since it varies among different contributors, leading to data sources that are quite heterogeneous (Goodchild 2007; Imran et al. 2015). For example, as Flickr provides users with manual geotagging, the photograph might be geotagged in a place by one user that differs from where it was taken in practice. In this case, the results would not be able to accurately identify areas of interest. In addition, social media use is self-selecting; users may not necessarily be representative of everyone who lives in or visits a city. For example, the primary user age group of Flickr is between 35 and 44 (Kahootz Media 2018). In addition, usage of the Flickr service is rather uneven, with more active users contributing to a larger number of photographs (Davies 2016; Hollenstein and Purves 2010). Such issues imply that research results may be focused on a particular segment of the population, and thus warrants caution when drawing conclusions. However, the degree of penetration and popularity of Flickr is such that we argue our results are still meaningful and can help us better understand urban dynamics from the perspective of people who experience the city through these lenses.

### 3.2 Data pre-processing

Data obtained directly from the API was preprocessed before analysis in two main stages: (1) subdividing the data set, and (2) eliminating noise.

A visualisation of the spatial distribution of the downloaded geotagged photograph locations in London can be seen in Fig. 1a, where kernel density estimation (O’Sullivan and Unwin 2014) is applied. The darker the red the higher the density, thus implying that more photographs were taken in central London relative to



**Fig. 1** **a** Distribution of the spatial density of Flickr photographs in London from 2013 to 2015 using a kernel density visualisation and **b** relationship between the number of Flickr photographs taken by each user and the number of unique users in London (2013–2015)

peripheral areas. Indeed, 73.5% of our Flickr data are located within the Inner London<sup>2</sup> definition. Thus, the study extent was narrowed to that of Inner London. In terms of the time dimension, our interest is set on the seasonal variability at the monthly level. Thus, we further divided the data into 36 monthly slices covering the periods between the first and last day of each month.

Next, we needed to identify erroneous or noisy records. First, we considered those cases where a user uploaded a few photographs at identical geographic locations (i.e. at least two photographs geocoded with the same longitude and latitude by the same user) in a month. Many photographs at the same longitude and latitude (given the recorded degree of precision of the coordinates) are quite unlikely, and as such are classified as erroneous in terms of the location attribute, perhaps as a result of faulty hardware. To remove this effect, only one record for each of these users was maintained. A similar case arises when a user takes multiple photographs in the same second of one day at different places; we also removed these cases. More importantly, Fig. 1b shows that a small group of users contributes a large proportion of the photographs. These are known as ‘active users’ by Hollenstein and Purves (2010) and Hu et al. (2015). The figure shows that a single user may upload hundreds of photographs in a year, while most users only upload dozens. In our definition of a UAOI, we argue that these entities should be agreed upon by many people, and the dominance of an active user may lead to bias in extracted UAOIs. An overemphasis on contributed content from any one user or subset of users (and their associated interests) will also influence the generality of the definitions of the UAOIs. To reduce the impact that such active users may have on shaping the outcomes of the analysis, we implemented a further set of cleaning routines that reduced the proportion of photographs from active users, by keeping only one photograph for each user based on tags used and the time when the photograph was taken. Specifically, if a user took several photographs in a minute but with the same tags, only one photograph was retained. The rationale for this approach was to remove photographs within a limited spatial extent on the hypothesis that people’s average walking speed is 5 km/h (Onaverage 2017). On this basis, the maximum walking distance within a minute is approximately 83 m. Within this distance, only a single user’s photograph that has the same text is retained. The specific data pre-processing steps are summarised in Table 2, showing how many photographs and users are removed following each step. After this process, an average number of 12,228 photographs and 2275 unique users remained in each month.

## 4 Methodological framework

In the following section, we present a systematic framework designed to extract and map the evolution of UAOIs from the subset of geotagged Flickr photographs outlined in the previous section. Our methodology consists of two main parts: cluster detection and boundary delineation.

---

<sup>2</sup> The Inner London definition comprises a series of centrally located London Boroughs within the Greater London Authority Extent (Mayor of London 2017).

**Table 2** The number of photographs and users at different stages of data pre-processing

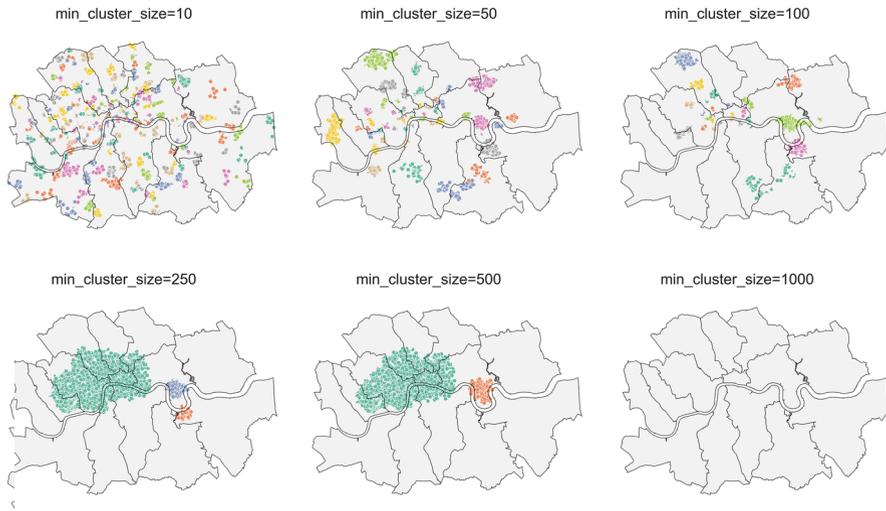
(1) Subdividing the data set	Total number of photos			Total number of unique users		
	Raw data collected in bounding box	1,579,694			39,531	
Data within Inner London	1,162,891			34,700		
(2) Eliminating noise	Monthly number of photos			Monthly number of unique users		
	Mean	Median	SD	Mean	Median	SD
Subdivide data into 36 time slices	32,221	32,991	5070.98	2287	2248	402.16
Remove photographs geocoded with identical coordinates from one user	16,116	15,899	2618.94	2287	2248	402.16
Remove systematic outliers	15,493	15,595	2342.96	2275	2244	393.02
Remove dominance of active users	12,228	11,913	1794.09	2265	2233	392.69

#### 4.1 Extracting urban areas of interest by the hierarchical density-based spatial clustering for applications with noise algorithm

We define UAOIs as those areas where multiple Flickr users have gathered and taken large numbers of spatiotemporally clustered photographs, reflecting a consensual view that some aspect of the urban environment is of interest. The extraction of such areas can be understood as a clustering problem, in particular, as one that has the aim of identifying robust, non-overlapping and dense concentrations of points. Following recent advances in the literature, we selected a density-based method. The advantages of such an approach are that they can produce results without pre-specification of cluster frequency and are robust to arbitrary shapes and the presence of outliers/noise deviating away from the main spatial distribution (Hans-Peter et al. 2011).

We applied the HDBSCAN (hierarchical density-based spatial clustering for applications with noise; Campello et al. 2013) as our clustering method as this overcomes several of the major drawbacks of other density-based algorithms. Contrary to more traditional algorithms, there is only one parameter to tune in HDBSCAN, with the other key parameter in the original DBSCAN implementation, i.e. the minimum cluster size (MinPts), being endogenously determined by the method. This approach represents a step forward in the direction of more robust, automated and data-driven techniques for the delineation of UAOIs. McInnes et al. (2017) describe the HDBSCAN process as comprising five steps:

1. Transform the space based on the estimates of density by defining a ‘mutual reachability’ distance, which is a new distance metric between points;
2. Build a minimum spanning tree to implement single-linkage clustering, which is a core feature of this algorithm;

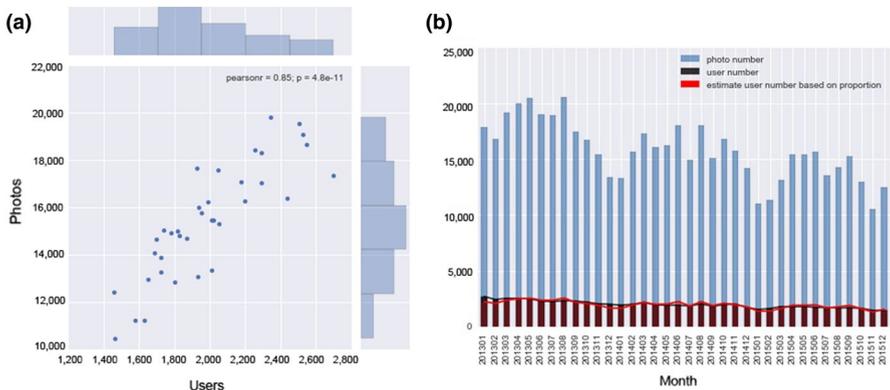


**Fig. 2** Different urban areas of interest extracted by different minimum cluster size (`min_cluster_size`) values in one month. Colours indicate the location of different clusters (colour figure online)

3. Construct a cluster hierarchy of connected components by iteratively sorting the edges of the tree by distance in an increasing order. The result can be viewed as a dendrogram that shows where robust single-linkage stops;
4. Condense the cluster hierarchy shown in the dendrogram into a smaller tree by attaching more data to each node;
5. Extract clusters that persist and are robust from the condensed tree.

Operationally, various epsilon values are generated automatically by the different density levels resulting from the single-linkage hierarchy, which allows HDBSCAN to find clusters of various densities. Also, it ensures improvements over OPTICS and DBSCAN by providing a clustering hierarchy, where a simplified tree of the most significant clusters (i.e. maximised stability) can be easily extracted.

When using HDBSCAN, the only parameter to specify is the minimum cluster size (`mclSize`), representing the minimum number of points (i.e. Flickr photographs) required for a UAOI to exist. In order to select an appropriate `mclSize`, we extensively explored the sensitivity of the final solution to changes in the parameter. A few representative thresholds, from 10 to 1000, were set as the minimum cluster size (`mclSize`) parameter, which were applied in all time slots. Figure 2 presents example outputs from this sensitivity analysis. We can see that if the `mclSize` is small (e.g. 10 or 50), more UAOIs are identified but there are also greater numbers of points labelled as noise (i.e. not part of any clusters); if the `mclSize` is larger (e.g. 500 or 1000), more robust results emerge, although clusters are significantly larger, causing potentially interesting but smaller areas to be missed. Furthermore, due to the number of Flickr photographs and users varying between months, it could be argued as being inappropriate to assign an absolute value for all time sequences. To handle these issues, values of 1–4% of



**Fig. 3** Exploring the relationship between Flickr photographs and users to ensure each urban area of interest contains multiple users. **a** Correlation analysis and **b** estimated proportion

the Flickr photographs in each month were assigned to *mcISize* across different iterations as discussed previously in order to produce appropriate frequencies of groups that fit the definition of a UAOI. After multiple experimental results, 1% of Flickr photographs in each month were used as the value for the minimum cluster size parameter, ensuring a higher number of UAOIs but also being cognisant of smaller clusters that may be of relevance.

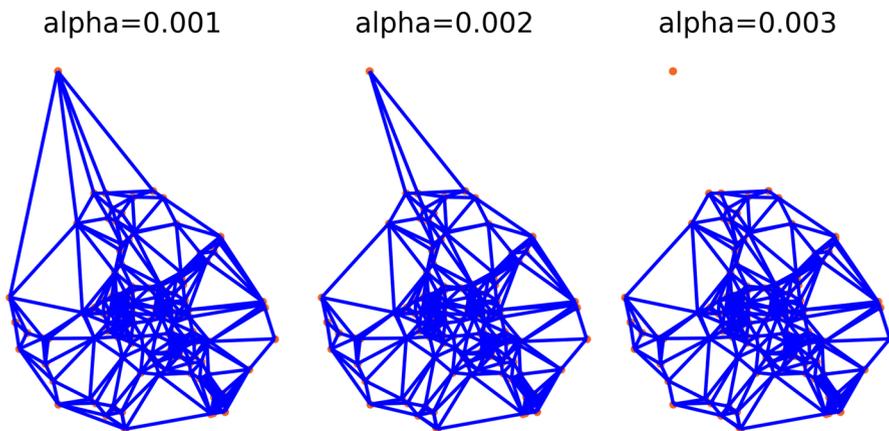
As UAOIs should be formed through the collective actions of multiple users within each specific time slice, the 1% parameter selection does not ensure that a set number of Flickr users are captured in each UAUI. As such, it was then necessary to verify the practical significance of the extracted UAOIs. An intuitive approach is to examine the relationship between the number of Flickr photographs and the number of users in each month. If they are correlated, then we can estimate the number of Flickr users by the number of photographs per month. Specifically, the scatter plot in Fig. 3a shows that there is a high positive correlation between the two variables, with a Pearson coefficient of 0.85, implying that as the number of photographs increases, so too does the number of users in a given UAUI. A linear regression model was then fitted using these two variables so that the user number could be estimated based on the number of photographs in each month. The resulting R-Squared was 0.725 with a *p* value for the coefficient value below 0.05, implying that the model is statistically significant and 72.5% of the variation in photograph numbers could be explained by the model. Figure 3b is a graph presenting the number of photographs, users, and the calculated user number in various time sequences. The red line fluctuates slightly around the black line, meaning that the 1% photograph number as the HDBSCAN parameter value can be interpreted as having at least 1% of users in each UAUI, which satisfies our definition of a UAUI. Therefore, we adopt these clustering results for the next stage of the analysis, which turns clusters of points into polygon boundaries.

## 4.2 Constructing a perceptual boundary to enclose the extracted urban areas of interest

The clusters from the method described above are represented as a group of points. However, within this study, we are interested in extracting largely non-overlapping UAIOs that refer to an area within a specific border. In other words, we are interested in identifying polygons rather than sets of points. The reason behind this procedure is twofold. First, as mentioned when introducing the concept, a UAIO was defined as a section of the city with an extraordinarily large density of images. Under this definition, two overlapping UAIOs would simply be merged into one. Secondly, our focus is to quantify spatiotemporal changes in the shape and extent of these polygons. In this context, even though a UAIO is identified with fixed borders at each point in time, its definition over time is much vaguer and is allowed to change, evolve and morph in line with changes to its underlying structure.

As such, the next step involves the construction of boundaries that enclose all geotagged images identified as part of a UAIO cluster. To delineate these shapes, we adopted a variant of the concave hull algorithm: the alpha shapes (Edelsbrunner et al. 1983). Alpha shapes are a widely used, robustly tested algorithm that create a tighter boundary as compared to the traditional convex hull method, which may produce large empty areas that do not belong to the original point data set (Akdag et al. 2014).

An alpha shape, which is a geometric concept, is a linear approximation of an original shape. It is a generalisation of the convex hull, and a subgraph of the Delaunay Triangulation (Edelsbrunner et al. 1983). It establishes a connection between each point and nearby points and removes the furthest triangles that are away from their neighbours. In this context,  $\alpha$  is a parameter that controls the desired level of detail, ranging from the standard “crude” convex hull ( $\alpha = \infty$ ) to the set of points itself ( $\alpha = 0$ , Da 2018). The algorithm first computes a Delaunay triangulation of the set of points ( $S$ ) and for each Delaunay edge, it computes the values  $\alpha$ -min ( $e$ ) and  $\alpha$ -max ( $e$ ). Next, for each edge  $e$ , if  $\alpha$ -min ( $e$ )  $\leq \alpha \leq \alpha$ -max ( $e$ ), the edge is kept in the  $\alpha$ -shape of  $S$ . We have tailored this general method to our application by developing a technique to find the most appropriate alpha value for each cluster. Like the parameter selection in HDBSCAN clustering, an absolute alpha value for all point clusters would not be suitable in that some areas would contain more empty areas in the range from 0.001 to 0.005. We then identified the first case where a single point was excluded from the main polygon and selected the previous value of alpha. This strategy resulted in the tightest polygon that still contained every point in the cluster. As an illustration, Fig. 4 represents three examples of different UAIOs produced with varying alpha values. In this case, 0.003 excludes a point (which in the original algorithm is still linked through an edge, but not an area), and 0.001 implies too sparse a solution compared to 0.002, which allows a tighter shape that still includes all points in the cluster within the same polygon. Hence, the value selected for this case is 0.002.



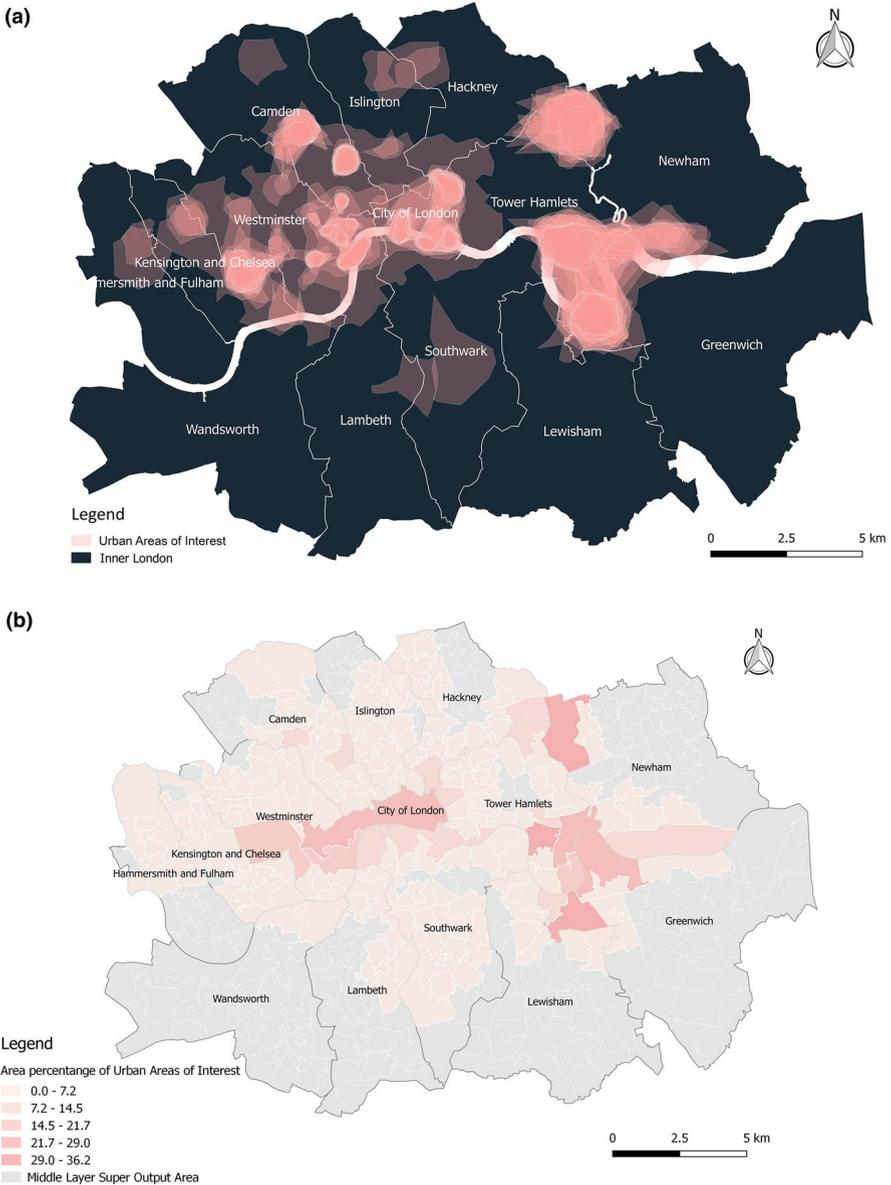
**Fig. 4** An example of one urban area of interest that changes with different alpha values for one month of data

## 5 Results

After applying the method described above to the geotagged Flickr photograph data set of Inner London from 2013 to 2015, the UAOIs were extracted for the 36 monthly slices. The spatiotemporal characteristics of the results are presented in this section.

We begin from a purely spatial perspective, “compressing” the temporal dimension. This approach allowed us to gain an idea of the stability of different parts of the city in being identified as UAOIs. Figure 5 presents each UAOI together in a single map. Figure 5a is produced by overlaying all UAOIs from different time sequences with a large degree of transparency to show the spatial distribution of the more stable UAOIs. Areas in darker pink are thus consistently identified as being of interest during the 3-year period, including: Trafalgar Square, St. Pancras International and tube station, King’s Cross, Jubilee Gardens, Westminster Pier, Borough Market, Millennium Bridge, Tower Bridge, the Canary Wharf financial centre, and the museums located on Museum Lane. These represent popular tourist attractions, cultural venues, business centres and locations with intense traffic.

Figure 5b is generated by aggregating the results of our analysis at the administrative boundary level, i.e. the middle-layer super output area (MSOA). MSOAs are designed to improve the reporting of small area (neighbourhood) statistics and are built from a hierarchy of output areas (OAs; Office for National Statistics 2018). These areas are intermediate in size between output areas and local authorities. Our intention with Fig. 5b was to transfer the extent to which a given part of the city belongs to a UAOI into a fixed geography that can be analysed over time. The map displays the total area identified as a UAOI in each MSOA over the entire period considered. The map effectively represents those small-scale areas that are more popular, shifting the attention from the organically evolving shapes of UAOIs to the more stable boundaries of MSOAs. The overall pattern displayed is similar to that



**Fig. 5** **a** All urban areas of interest extracted in inner London from 2013 to 2015 showing the most stable and popular spatial zones and **b** the overall spatial distribution of the total area of the urban areas of interest in each middle-layer super output area

in Fig. 5a, showing higher values in the northwest of Newham, the border of Tower Hamlets and Greenwich, the City of London and the middle of Westminster borough, implying a higher degree of attention in these districts.

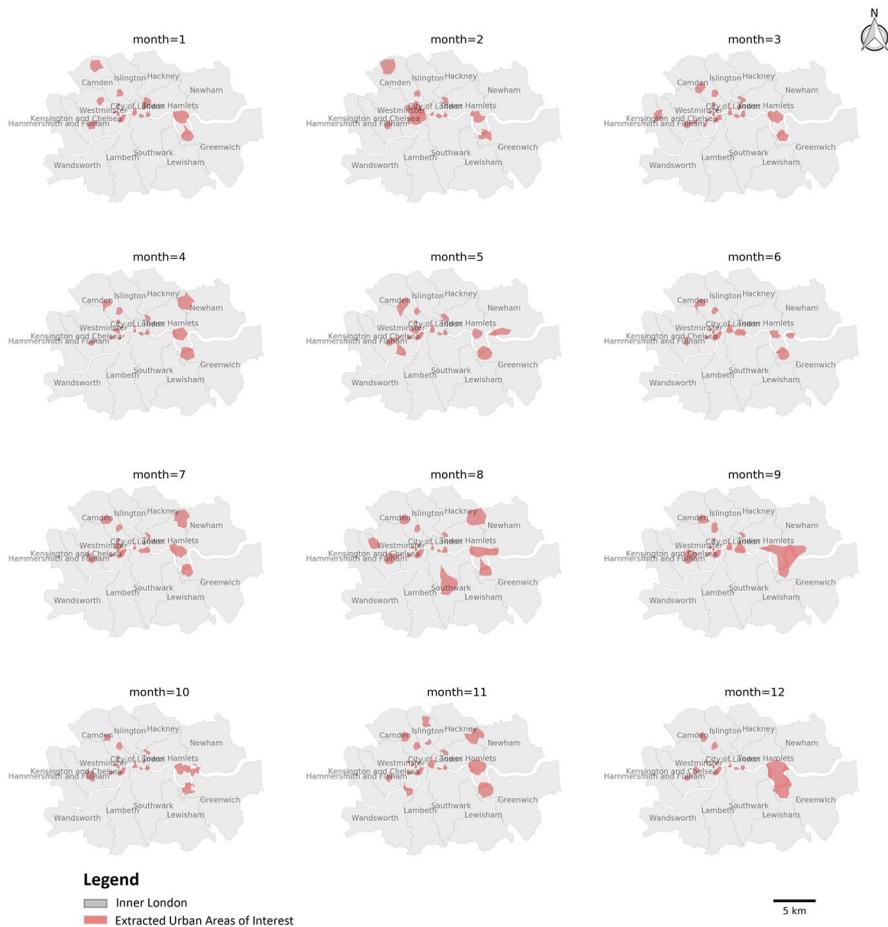
Although by the nature of the analysis and the source of data employed, it is very hard to carry out a formal validation of the results, the patterns displayed in Fig. 5a, b are well aligned with established knowledge from the literature. Both maps result from the interaction between the urban built environment and human behaviour and highlight popular areas generally covering business centres, public entertainment (theatre, Art Centre and Sports Centre) and food markets, as well as open spaces. They also illustrate that people are more likely to take photographs in those regions where most of the significant landmarks and unique buildings are located. A good example is the City of London, which contains a historical centre with historical buildings as well as modern skyscrapers, and serves as a central business district. We can also see that the districts on the border of Tower Hamlets and Hackney are not always identified as part of a UAOI, which suggests that the degree of popularity of these districts is influenced by different factors and may vary seasonally.

The temporal nature of UAOIs is explored in Fig. 6, which shows how their extent changes during a single year (i.e. 2013). We can see that some UAOIs emerged and disappear suddenly in the span of 1 or 2 months, which indicates that there is a high probability that large-scale but temporary events took place in these areas. For example, the UAOI extracted in the north of Camden existed only in January and February and then disappeared during the following months. This is likely caused by the first snowfall in London in January 2013, as Hampstead Heath is known as a good place for people to enjoy snow by sledging, activities that are usually recorded in photographs. This event was reported in multiple media (Emms 2013; Pettitt 2013).

Although useful, it is difficult to scale the spatiotemporal variation in Fig. 6. Every additional month involves a full map, and comparing a large number of maps at the same time carries a large cognitive load. To be able to extend the analysis and consider the entire period of 3 years at a fine temporal resolution, we created area profiles for stable geographical entities. We designed this approach to avoid directly examining and comparing the shape of each UAOI over time, as it is difficult and unintuitive to track and follow change with such an approach. Because of their organic and rapidly evolving nature, their shape and extent may vary significantly over time. This makes consistent temporal analysis complicated if the original shapes are to be used. For this reason, we returned to the MSOAs. Area profiles are a series of time plots that display, for every MSOA, the percentage of the area that is considered part of a UAOI in a given month. These figures are able to intuitively summarise the degree of participation of a given MSOA in UAOIs, as well as their evolution over the period considered, jointly capturing space and time in a single figure. To put this profile into context, the time plot is complemented with a map that shows the location of the area considered.

Figure 7 shows the UAOI profiles of three MSOAs with distinct characteristics throughout the 3 years from January 2013 to December 2015. These spatiotemporal profiles can thus help stakeholders better understand the dynamic characteristics of these districts when, for example, allocating resources more effectively, or enhance their understanding of the seasonal interest in specific geographic areas of the city.

The first profile corresponds to an area in Westminster. The profile clearly shows a seasonal evolution, oscillating around 15–20%, with higher percentages in warmer



**Fig. 6** The spatiotemporal evolution of urban areas of interest in 2013

months (June, July and August), and lower participation in UAOIs in colder months. In addition, there are also three outliers corresponding to February 2013, and January and February 2015, which display a larger share of the area being part of a UAOI. In particular, the 2015 outliers reach the full extent of the MSOA. It is hard to tell why these occurred, and an in-depth exploration of each of these warrants further research (e.g. semantic analysis or image recognition), which is beyond the scope of this paper. However, what they help to highlight is the ability of the profile to make these patterns explicit and alert the analyst about their existence in a way that traditional maps do not. The ability is even clearer if we consider the profile of the area in Tower Hamlets. In this case, the seasonal variation is more pronounced, moving from about 20% to the entire coverage of the MSOA. These spikes are not necessarily outliers, as they occur in each of the 3 years considered during the warmer months. The only one that could be considered an anomaly is that of March 2014, which took place at a time outside the summer period. Equally, the



**Fig. 7** Spatiotemporal profiles for urban areas of interest based on middle-layer super output layer geographic areas

MSOA was not part of any UAIO during November and December of 2015 which, compared to the previous years, was expected. Again, these patterns warrant further research to explore the drivers behind them, but the role of the profile in highlighting them is clear. Finally, the third panel in Fig. 7 shows a different type of area. The Newham example displays several months in which the area is not part of any UAIO. However, the spring and summer months see it consistently having around a third of its extension within an identified UAIO. This pattern implies that the popularity of this district is significantly influenced by season and its role in the overall hierarchy is less prominent than that of the other two areas considered here.

## 6 Discussion and conclusions

This paper provides insight into several questions relevant for research concerned with VGI as a means of better understanding urban environments. We propose a framework to extract UAIO boundaries from geotagged image data, and use them to build spatiotemporal profiles of areas. When compared to existing literature, our approach is distinct in two key dimensions. First, we introduce the use of the recent HDBSCAN clustering algorithm, which we show improves on the results of other commonly used density-based algorithms employed in previous studies (Kisilevich et al. 2010a; Hollenstein and Purves 2010; Li et al. 2013; Lee et al. 2014; Gao et al. 2017). Second, our approach is significantly more detailed in terms of temporal resolution, which allows us to characterise areas based on their seasonal profiles. This again brings a new perspective to previous approaches (e.g. Andrienko and Andrienko 2013; Hu et al. 2015), which focus on coarser temporal scales.

The results on the spatial dimension of our analysis suggest that the urban environment influences human activity, shaping the attention of people and attracting them to areas where many unique buildings and important landmarks are located. Conversely, the temporal aspect of our results reflects how human activity evolves and shapes the use of the urban environment. Putting these two together, our spatiotemporal profiles visualise how the popularity of certain regions is influenced by factors such as time of the year and season, and also make visually explicit how popularity levels differ across areas. This approach is distinct from related works that use VGI to study UAOIs, such as Hu et al. (2015), in that our perspective is more granular and thus allows us to uncover qualitatively different types of dynamics. Spatially, we are focused on the internal dynamics of urban environments and in comparing areas within the same city. Temporally, we use higher resolution to consider seasonal changes, rather than longer-term evolution.

The methods and results presented in this paper are of interest for several fields and domains. For example, it can help urban planners to develop better strategies related to tourism planning. If certain tourist attractions showed a seasonal pattern according to the spatiotemporal profiles produced in this study, urban planners could allocate resources for tourism more efficiently. Local authorities may also be interested in those UAOIs that are the most stable and have a larger area throughout the year for purposes such as police patrol and traffic monitoring. The results can also be used by researchers and practitioners as an additional geographic layer to understand the use of the urban built environment. Furthermore, part of the relevance of our contribution lies in the fact that it can be deployed using data that are available in near real-time. Unlike more traditional data sources, geotagged images are constantly added to services such as Flickr, thus providing an opportunity to study the evolution of UAOIs not only retrospectively but as they evolve over time. This holds distinct value for practitioners such as urban planners and policy makers.

There are several avenues towards which the work presented in this paper could be extended. Although the data set used here is extracted from Flickr, geotagged images from other websites could be used. Different platforms provide slightly different services that attract different populations (Lazer and Radford 2017). Incorporating different sources would thus likely improve the coverage of the analysis and provide a novel comparison of the inherent biases of each platform. Additionally, our current focus has been on the spatial and temporal aspects of the images. A promising further avenue for research is to include information in the analysis other than spatiotemporal stamps such as, for example, the text included in tags, or the images themselves. The former would expand existing work on semantic ontologies (Kisilevich et al. 2010a; Lee et al. 2014), while the latter would complement recent advances on deep learning that aim at extracting features from images (Krizhevsky et al. 2017; Zhou et al. 2017; Redmon and Farhadi 2017). Finally, this analysis could also be further extended by considering the socioeconomic characteristics of Flickr users, seeking to establish a link between, e.g. Flickr metadata and census data. These applications, although beyond the scope of this present paper, warrant future attention by researchers.

**Open Access** This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

## References

- Akdag F, Eick CF, Chen G (2014) Creating polygon models for spatial clusters. In: Andreasen T, Christiansen H, Cubero JC, Raś ZW (eds) Foundations of intelligent systems. ISMIS 2014. Lecture notes in computer science, vol 8502. Springer, Cham. [https://doi.org/10.1007/978-3-319-08326-1\\_50](https://doi.org/10.1007/978-3-319-08326-1_50)
- Andrienko G, Andrienko N (2013) Thematic patterns in georeferenced tweets through space-time visual analytics. *Comput Sci Eng* 15(3):72–82. <https://doi.org/10.1109/MCSE.2013.70>
- Ankerst M, Breunig M, Kriegel H, Sander J (1999) OPTICS: ordering points to identify the clustering structure. *ACM SIGMOD Rec* 28(2):49–60. <https://doi.org/10.1145/304182.304187>
- Antoniou V, Fonte C, See L, Jokar Arsanjani J, Lupea F, Minghini M, Foody G, Fritz S (2016) Investigating the feasibility of photographs as sources of land cover data. *Int J Geo-Inf* 5(5):64. <https://doi.org/10.3390/ijgi5050064>
- Arribas-Bel D (2014) Accidental, open and everywhere: emerging data sources for the understanding of cities. *Appl Geogr* 49:45–53. <https://doi.org/10.1016/j.apgeog.2013.09.012>
- Batty M (2017) The age of the smart city. <https://doi.org/10.13140/RG.2.2.34783.87204>
- Campello RJGB, Moulavi D, Sander J (2013) Density-based clustering based on hierarchical density estimates. In: Pei J, Tseng VS, Cao L, Motoda H, Xu G (eds) Advances in knowledge discovery and data mining. PAKDD 2013. Lecture notes in computer science, vol 7819. Springer, Berlin. [https://doi.org/10.1007/978-3-642-37456-2\\_14](https://doi.org/10.1007/978-3-642-37456-2_14)
- Catt RD (2009) 100,000,000 geotagged photos (plus). <http://code.flickr.net/2009/02/04/100000000-geotagged-photos-plus/>. Accessed 30 May 2017
- Chourabi H, Nam T, Walker S, Gil-Garcia JR, Mellouli S, Nahon K, Pardo TA, Scholl HJ (2012). Understanding Smart City initiatives: an integrative framework. In: 45th Hawaii international conference on system sciences, pp 2289–2297, 4–7 January Maui, HI. <https://doi.org/10.1109/HICSS.2012.615>
- Crandall DJ, Backstrom L, Huttenlocher D, Kleinberg J (2009) Mapping the world's photos. In: Proceedings of the 18th international conference on World wide web, ACM, New York, NY, USA, pp 761–770. <https://doi.org/10.1145/1526709.1526812>
- Crooks A, Croitoru A, Jenkins A, Mahabir R, Agouris P, Stefanidis A (2016) User-generated big data and urban morphology. *Built Environ* 42(3):396–414. <https://doi.org/10.2148/benv.42.3.396>
- Da TKF (2018) 2D alpha shapes. In: CGAL user and reference manual, 4.13 edn. CGAL Editorial Board. <https://doc.cgal.org/latest/Manual/packages.html#PkgAlphaShape2Summary>. Accessed 1 Nov 2018
- Davies J (2016) The virtual and the physical. <http://www.nesta.org.uk/blog/virtual-and-physical/>. Accessed 24 Apr 2017
- Edelsbrunner H, Kirkpatrick DG, Seidel R (1983) On the shape of a set of points in the plane. *IEEE Trans Inf Theory* 29(4):551–559. <https://doi.org/10.1109/TIT.1983.1056714>
- Emms S (2013) A proper snow day on hampstead heath—in pictures. Kentishtowner. <https://www.kentishtowner.co.uk/2013/01/20/snow-day-on-hampstead-heath-2013-for-real/>. Accessed 9 Sept 2017
- Ester M, Kriegel HP, Sander J, Xu X (1996) A density-based algorithm for discovering clusters in large spatial databases with noise. In: KDD'96 proceedings of the second international conference on knowledge discovery and data mining, pp 226–231
- Foursquare for Developer (2017) Policies. <https://developer.foursquare.com/overview/community>. Accessed 18 May 2017
- Gandy M (2006) Urban nature and the ecological imaginary. In: Swyngedouw E, Kaika M, Heynen N (eds) In the nature of cities: urban political ecology and the politics of urban metabolism. Routledge, London, pp 63–74
- Gao S, Janowicz K, Montello DR, Hu Y, Yang JA, McKenzie G, Yan B (2017) A data-synthesis-driven method for detecting and extracting vague cognitive regions. *Int J Geogr Inf Sci* 31(6):1245–1271. <https://doi.org/10.1080/13658816.2016.1273357>
- Goodchild MF (2007) Citizens as sensors: the world of volunteered geography. *GeoJournal* 69(4):211–221. <https://doi.org/10.1002/9780470979587.ch48>

- Hans-Peter K, Peer K, Jörg S, Arthur Z (2011) Density-based clustering. *WIREs Data Min Knowl Discov*. <https://doi.org/10.1002/widm.30>
- Harrison C, Eckman B, Hamilton R, Hartswick P, Kalagnanam J, Paraszczak J, Williams P (2010) Foundations for smarter cities. *IBM J Res Dev* 54(4):1–16. <https://doi.org/10.1147/JRD.2010.2048257>
- Hollenstein L, Purves RS (2010) Exploring place through user-generated content: using Flickr to describe city cores. *J Spat Inf Sci* 1(1):21–48. <https://doi.org/10.5311/JOSIS.2010.1.3>
- Hu Y, Gao S, Janowicz K, Yu B, Li W, Prasad S (2015) Extracting and understanding urban areas of interest using geotagged photos. *Comput Environ Urban Syst* 54:240–254. <https://doi.org/10.1016/j.compenvurbsys.2015.09.001>
- Imran M, Castillo C, Diaz F, Vieweg S (2015) Processing social media messages in mass emergency: a survey. *ACM Comput Surv* 47(4):1–38. <https://doi.org/10.1145/2771588>
- Jiang S, Alves A, Rodrigues F, Ferreira J, Pereira FC (2015) Mining point-of-interest data from social networks for urban land use classification and disaggregation. *Comput Environ Urban Syst* 53:36–46. <https://doi.org/10.1016/j.compenvurbsys.2014.12.001>
- Kahootz Media (2018) Social Media Demographics 2017—Who's Using What? <http://www.kahootzmedia.com/social-media-demographics-2017-whos-using/>. Accessed 13 Apr 2018
- Kisilevich S, Krstajic M, Keim D, Andrienko N, Andrienko G (2010) Event-based analysis of people's activities and behaviour using Flickr and Panoramio geotagged photo collections. In: *Information Visualisation (IV), 2010 14th international conference*, pp 289–296. IEEE. <https://dx.doi.org/10.1109/IV.2010.94>
- Kisilevich S, Mansmann F, Bak P, Keim D, Tchaikin A (2010) Where would you go on your next vacation? A framework for visual exploration of attractive places. In: *Advanced geographic information systems, applications, and services, 2010 Second international conference*, pp 21–26. IEEE. <https://doi.org/10.1109/GEOProcessing.2010.11>
- Krizhevsky A, Sutskever I, Hinton G (2017) ImageNet classification with deep convolutional neural networks. *Commun ACM* 60(6):84–90. <https://doi.org/10.1145/3065386>
- Lansley G, Longley PA (2016) The geography of twitter topics in London. *Comput Environ Urban Syst* 58:85–96. <https://doi.org/10.1016/j.compenvurbsys.2016.04.002>
- Lazer D, Radford J (2017) Data ex machina: introduction to big data. *Ann Rev Sociol* 43:19–39. <https://doi.org/10.1146/annurev-soc-060116-053457>
- Lee I, Cai G, Lee K (2014) Exploration of geo-tagged photos through data mining approaches. *Expert Syst Appl* 41(2):397–405. <https://doi.org/10.1016/j.eswa.2013.07.065>
- Li L, Goodchild MF (2012) Constructing places from spatial footprints. In: *Proceedings of ACM SIGSPATIAL GEOCROWD'12, Redondo Beach, CA, USA*. <https://doi.org/10.1145/2442952.2442956>
- Li L, Goodchild MF, Xu B (2013) Spatial, temporal, and socioeconomic patterns in the use of Twitter and Flickr. *Cartogr Geogr Inf Sci* 40(2):61–77. <https://doi.org/10.1080/15230406.2013.777139>
- Lloyd A, Cheshire J (2017) Deriving retail centre locations and catchments from geo-tagged Twitter data. *Comput Environ Urban Syst* 61:108–118. <https://doi.org/10.1016/j.compenvurbsys.2016.09.006>
- Mayor of London (2017) Policy 2.9 Inner London. <https://www.london.gov.uk/what-we-do/planning/london-plan/current-london-plan/london-plan-chapter-two-londons-places/policy-29/>. Accessed 3 May 2017
- McInnes L, Healy J, Astels S (2017) hdbscan: hierarchical density-based clustering. *J Open Source Softw Open J*. <https://doi.org/10.21105/joss.00205>
- Michel F (2017) How many public photos are uploaded to Flickr every day, month, year? <https://www.flickr.com/photos/franckmichel/6855169886/>. Accessed 5 Aug 2017
- Office for National Statistics (2018) Census geography. <https://www.ons.gov.uk/methodology/geography/ukgeographies/censusgeography#super-output-area-soa>. Accessed 3 Sept 2018
- Onaverage (2017) Average walking speed. <http://www.onaverage.co.uk/speed-averages/average-walking-speed/>. Accessed 26 Aug 2017
- O'sullivan D, Unwin D (2014) *Geographic information analysis*. Wiley, Hoboken
- Papadopoulos S, Zigkolis C, Kompatsiaris Y, Vakali A (2011) Cluster-based landmark and event detection for tagged photo collections. *IEEE Multimedia* 18(1):52–63. <https://doi.org/10.1109/MMUL.2010.68>
- Pettitt J (2013) Picture gallery: fun in the snow on Hampstead Heath. *Hampstead Highgate Express*. <http://www.hamhigh.co.uk/news/picture-gallery-fun-in-the-snow-on-hampstead-heath-1-1801567>. Accessed 15 Mar 2018
- Rattenbury T, Good N, Naaman M (2007) Towards automatic extraction of event and place semantics from Flickr tags. In: *Proceedings of the 30th annual international ACM SIGIR conference on*

- research and development in information retrieval, pp 103–110. ACM. <https://doi.org/10.1145/1277741.1277762>
- Redmon J, Farhadi A (2017). YOLO9000: better, faster, stronger. In: CVPR
- Shi K, Yu B, Huang Y, Hu Y, Yin B, Chen Z, Chen L, Wu J (2014) Evaluating the ability of NPP-VIIRS nighttime light data to estimate the gross domestic product and the electric power consumption of China at multiple scales: a comparison with DMSP-OLS data. *Remote Sens* 6(2):1705–1724. <https://doi.org/10.3390/rs6021705>
- Smith, C. (2016). 14 Interesting Flickr Stats. DMR. <http://expandedramblings.com/index.php/flickr-stats/>. Accessed 9 May 2017
- Stüvel S (2016) Python Flickr API kit—Stüvel Stuff. <https://stuvel.eu/flickrapi>. Accessed 5 Apr 2017
- Sun Y, Fan H, Bakillah M, Zipf A (2015) Road-based travel recommendation using geo-tagged images. *Comput Environ Urban Syst* 53:110–122. <https://doi.org/10.1016/j.compenvurbsys.2013.07.006>
- Tasse D, Hong J (2014) Using social media data to understand cities. In: Proceedings of NSF workshop on big data and urban informatics
- Wang R, Strong D (1996) Beyond accuracy: what data quality means to data consumers. *J Manag Inf Syst* 12(4):5–33. <https://doi.org/10.1080/07421222.1996.11518099>
- Washburn D, Sindhu U, Balaouras S, Dines RA, Hayes NM, Nelson LE (2010) Helping CIOs understand “Smart City” initiatives: defining the smart city, its drivers, and the role of the CIO. Forrester Research, Cambridge
- Yanai K, Yaegashi K, Qiu B (2009). Detecting cultural differences using consumer-generated geotagged photos. In: Proceedings of the 2nd international workshop on location and the web, p 12. ACM. <https://doi.org/10.1145/1507136.1507148>
- Zheng YT, Zha ZJ, Chua TS (2012) Mining travel patterns from geotagged photos. *ACM Trans Intell Syst Technol* 3(3):56. <https://doi.org/10.1145/2168752.2168770>
- Zhou B, Lapedriza A, Torralba A, Oliva A (2017) Places: a 10 million image database for scene recognition. *IEEE Trans Pattern Anal Mach Intell*. <https://doi.org/10.1109/TPAMI.2017.2723009>