

## Streets of London: Using Flickr and OpenStreetMap to build an interactive image of the city



Azam Raha Bahrehdar<sup>a,\*</sup>, Benjamin Adams<sup>b</sup>, Ross S. Purves<sup>a</sup>

<sup>a</sup> Department of Geography, University of Zurich, Winterthurerstrasse 190, Zurich 8057, Switzerland

<sup>b</sup> Department of Computer Science and Software Engineering, University of Canterbury, Christchurch, New Zealand

### ARTICLE INFO

#### Keywords:

Streets  
Lynch  
Similarity  
User generated content  
Perception

### ABSTRACT

In his classic book “The Image of the City” Kevin Lynch used empirical work to show how different elements of the city were perceived: such as paths, landmarks, districts, edges, and nodes. Streets, by providing paths from which cities can be experienced, were argued to be one of the key elements of cities. Despite this long standing empirical basis, and the importance of Lynch’s model in policy associated areas such as planning, work with user generated content has largely ignored these ideas. In this paper, we address this gap, using streets to aggregate filtered user generated content related to more than 1 million images and 60,000 individuals and explore similarity between more than 3000 streets in London across three dimensions: user behaviour, time and semantics. To perform our study we used two different sources of user generated content: (1) a collection of metadata attached to Flickr images and (2) street network of London from OpenStreetMap. We first explore global patterns in the distinctiveness and spatial autocorrelation of similarity using our three dimensions, establishing that the semantic and user dimensions in particular allow us to explore the city in different ways. We then used a Processing tool to interactively explore individual patterns of similarity across these four dimensions simultaneously, presenting results here for four selected and contrasting locations in London. Before drilling into the data to interpret in more detail, the identified patterns demonstrate that streets are natural units capturing perception of cities not only as paths but also through the emergence of other elements of the city proposed by Lynch including districts, landmarks and edges. Our approach also demonstrates how user generated content can be captured, allowing bottom-up perception from citizens to flow into a representation.

### 1. Introduction

The tale of Dick Whittington tells the story of a poor country boy who, enticed by rumours of streets paved with gold, makes his way to London. On his arrival he finds a busy, dirty city where his senses are assailed by sounds, smells and sights which are very different from those he had imagined. If the fictional Dick Whittington were alive today, he might use social media to take pictures of London and document some of the things he saw. By analysing not only what he photographed, but also comparing it to what others described, we could perhaps have a way of characterising London. But presumably Dick Whittington’s descriptions would be rather different to those of London’s inhabitants, for whom the noise and bustle experienced by the country boy are simply background noise. And perhaps some locations, say a city market, would have distinct temporal signatures, reflecting how use of space varies over time. Other spaces might be preferred by locals, and not visited by Dick Whittington and other recent arrivals to

the city. How we might extract and use such information to better understand how cities are perceived, by whom, and when is the subject of this paper.

The first question that we must answer in such an endeavour is what are the parts which come together in our perception of a city? Lynch, in his seminal book “The Image of the City” argued that cities are perceived through five elements: paths, nodes, districts, edges and landmarks. Paths, he claimed, are “channels along which the observer ... moves” and included, importantly for our work, streets which were for many people the “predominant elements” in their image of the city (Lynch, 1960). This importance of paths or streets is widely recognised in urban planning – both in terms of their function in enabling mobility and as experienced public spaces (von Schnfeld & Bertolini, 2017). Districts were described by Lynch as “the relatively large city areas which the observer can mentally go inside of, and which have some common character.” Districts contain not only paths, but also salient landmarks in mental maps of the city. Nodes include locations linking

\* Corresponding author.

E-mail addresses: [azam.bahrehdar@geo.uzh.ch](mailto:azam.bahrehdar@geo.uzh.ch) (A.R. Bahrehdar), [benjamin.adams@canterbury.ac.nz](mailto:benjamin.adams@canterbury.ac.nz) (B. Adams), [ross.purves@geo.uzh.ch](mailto:ross.purves@geo.uzh.ch) (R.S. Purves).

<https://doi.org/10.1016/j.compenvurbsys.2020.101524>

Received 13 December 2019; Received in revised form 7 July 2020; Accepted 8 July 2020

0198-9715/© 2020 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

paths (such as squares) which may ease orientation. Edges are linear physical or cultural divisions which are often borders between districts, or barriers to paths. Lynch's model is only one possible way of partitioning a city, but its relative simplicity, its empirical grounding, and its prominence in urban planning make it attractive (e.g., [Hospers \(2010\)](#) and [Carmona, Heath, Oc, and Tiesdell \(2012\)](#)). The potential of streets as a way of immersing oneself in a virtual city—think of, for example, Google's Street View—is a further indicator that Lynch's paths are a logical group of elements through which to partition a city.

Given a set of objects to describe, a second key question is, with what? User generated content (UGC), in the form of images and their metadata have proven to be a tractable way of collecting perceptual information about locations which was previously the domain of empirical work (such as the questionnaires and interviews used by Lynch). The possibility of using such data at scale has spawned a vast literature exploring the utility of UGC as a data source. Particularly predominant with respect to the characterisation of cities has been research using Flickr, most likely due to relatively stable access, and the ability to query using a range of different dimensions including location. Early research demonstrated that clustering image locations and their descriptions could provide tractable ways of describing space ([Rattenbury & Naaman, 2009](#)), and showed how perceptual theory (for example, with respect to the elements tagged) was replicated in such data ([Rorissa, 2008](#); [Tversky & Hemenway, 1983](#)). The predominance of toponyms as tags led to a wide range of work focussing on the delineation of vague regions, analogous to the districts described by Lynch ([Hollenstein & Purves, 2010](#); [Hobel, Fogliarini, & Frank, 2016](#); [Gao et al., 2017](#)). Using Flickr tags as ways of characterising cities and popular landmarks was quickly exploited ([Crandall, Backstrom, Huttenlocher, & Kleinberg, 2009](#)), essentially capturing cities at two contrasting granularities—as aggregate entities and through individual, highly popular cultural attractions analogous to Lynch's landmarks. Further analysis of image tags revealed that they captured not only visually perceived elements, but also allowed inferences to be made about sounds and smells ([Quercia, Schifanella, Aiello, & McLean, 2015](#)) in the city. Through these, and other qualities, it was possible to map potential preferences throughout a city, and thus recommend, for example, beautiful paths [Quercia, Schifanella, & Aiello, 2014](#)). However, despite the emphasis Lynch places on paths as key contributors to a city's image, very few UGC studies have focused on data aggregated at the granularity of path-like features such as streets.

Identifying and characterising similar streets in a city has numerous applications. For example, it can be used, as suggested above, in route recommendation, or also more generally in recommender systems ([Huang, 2016](#); [Quercia et al., 2014](#)). Similar streets may help to identify meaningful units at different scales, such as the districts proposed by Lynch, and gaps in similarity may suggest potential edges, or barriers of relevance in planning and tourism. By developing computational methods which capture such properties, we can develop tools which might help bridge gaps between quantitative and qualitative methods, by allowing researchers to explore a large space through methods akin to what is known in the digital humanities as *macroreading* before zooming in to apply more qualitative methods to, for instance, compare streets which appear semantically similar based on UGC. In this paper we take a first step towards these aims, demonstrating how, using streets as a proxy for Lynch's paths, we can use image metadata to characterise and compare within a city, using London as an example.

We extend existing work by linking efforts to characterise cities using UGC with emerging computational approaches capturing Lynch's ideas. We focus on comparing paths, one of the most important elements identified by Lynch, with a further clear need coming from urban studies, and yet strangely neglected in many works focusing on UGC. We do so by considering three dimensions of paths: the users who visit them, the ways in which they describe them and the times at which they are visited.

Our contribution is thus threefold:

- We demonstrate how UGC can be linked to paths allowing us to create a computational version of how a city is perceived after appropriate data filtering.
- We show how paths can be compared and ranked according to three contrasting dimensions: their descriptions, the users who visit them and their temporal signatures.
- We explore how and why contrasting dimensions capture similarity by comparing and interpreting signatures.

## 2. Background

Describing cities, and capturing the properties which make particular places within cities more or less distinctive, is a key task if we are to effectively digitally represent cities ([Miller & Small, 1999](#)). Rallying calls to consider place in geographic information science (e.g., [Goodchild \(2011\)](#) and [Elwood, Goodchild, and Sui \(2013\)](#)) have focused on the need to better capture shared, bottom-up representations of place, which go beyond categorisations of space derived from traditional, authoritative sources of spatial data. In particular, arguments for place-based representations often advance the idea of better representing varied human experiences of a location [Adams & McKenzie, 2013](#); [Jenkins, Croitoru, Crooks, & Stefanidis, 2016](#)), moving from the purely physical (e.g., park benches and bus stops) to, for example, emotions and behaviours associated with places ([Hauthal & Burghardt, 2013](#); [Shelton, Poorthuis, Graham, & Zook, 2014](#); [van Weerdenburg, Scheider, Adams, Spierings, & van der Zee, 2019](#)). The advent of social media, particularly user generated content, has provided an opportunity to capture human cognitive notion of place. While work like that of [Montello \(2003\)](#) focussed on capturing areas of interest through interviews, ([Hu et al., 2015](#), p. 1) argued that such approaches are labor-intensive, time-consuming, and do not scale well.

A key reason for the emergence of research on computationally representing place can therefore be linked to the data avalanche referred to by [Miller \(2010\)](#) with respect to the production of fine-grained data on urban spaces, and in particular rich UGC contributed by many individuals containing not only spatial information but also related temporal and semantic content. UGC, in different forms, has been used by many authors to characterise different dimensions of cities. One of the most prominent examples of such data are georeferenced microblog entries in Twitter. However, we note that though these data are suitable for exploring broad scale patterns of, for example language use or segregation in cities ([Shelton et al., 2014](#)), they have shortcomings with respect to fine-grained analysis ([Lansley & Longley, 2016](#)). On the one hand, attempts to georeference the large proportion of Tweets not explicitly furnished with coordinates typically fail at fine resolutions except when matching to select sets of commercial points of interest ([Zheng, Han, & Sun, 2018](#)), and, on the other hand, the content of a georeferenced Tweets was not always relevant to location ([Hahmann, Purves, & Burghardt, 2014](#)). By contrast, image descriptions, uploaded to image sharing platforms, have a number of desirable properties. Firstly, early work demonstrated that image tags contained sufficient semantics to allow meaningful descriptions to be generated for locations (e.g., [Rattenbury and Naaman \(2009\)](#) and [Crandall et al. \(2009\)](#)). Secondly, image tags capture not only visually perceived elements, but also inherent qualities of places including affordances and perceptual properties ([Dunkel, 2015](#)). Thirdly, since one reason why users tag images is to make them findable, image tags often reflect basic levels ([Rorissa, 2008](#))—they use shared terms which are both informative and succinct. Fourthly, data quality is good, such that image metadata containing coordinates are both accurate and precise, with caveats as to whether the location of the photographer or the subject are captured ([Hollenstein & Purves, 2010](#); [Zielstra & Hochmair, 2013](#)), allowing extraction of spatial properties of individual landmarks ([Crandall et al., 2009](#)).

These data properties have led to multiple studies based around Flickr images, their locations and associated metadata including tags,

timestamps and unique user identifiers. Early work transferred concepts from traditional information retrieval, such as term frequency-inverse document frequency weighting, to derive salient and distinctive descriptions for spatial regions (Kennedy, Naaman, Ahern, Nair, & Rattenbury, 2007). By analysing the locations of Flickr images and their tags, Crandall et al. (2009) showed that landmarks from different global cities could be extracted, and also demonstrated how the importance of salient landmarks in characterising different cities varied. Flickr also quickly proved to be an excellent source of information allowing vague places and vernacular names to be mapped at the city scale (e.g., Hollenstein and Purves (2010)). Understanding which parts of cities were visited, and in which order, requires that trajectories be built from images taken by individual users (Girardin, Calabrese, Fiore, Ratti, & Blat, 2008). In large urban areas, simply distinguishing between ‘locals’ and ‘tourists’, based on the length of time an individual is present, proved to be a very effective way of describing use of space as shown by Eric Fischer (2013) in a set of impressive visualisations.<sup>1</sup> Straumann, Cöltekin, and Andrienko (2014) use temporal and user information to build trajectories and compare group behaviours and thus, they argue, explore narratives in the city. Feick and Robertson (2015) make two important observations—firstly, the semantics derived from georeferenced image tags is dependent on the scale of the analysis unit and, secondly, the distribution of images is strongly influenced by the street network (and open spaces). This second observation makes it all the more surprising in our view that most studies to date have linked urban properties derived from UGC with space ignoring the underlying street network. Indeed, even work focusing on deriving “beautiful, quiet, and happy routes in the city” used a grid to characterise locations based on terms extracted from Flickr data and associated with, for example positive and negative emotions (Quercia et al., 2014). Finally, we note that though many studies have characterised and compared regions or grid cells using UGC (e.g., Derungs and Purves (2016) and Gao et al. (2017)) a detailed exploration of the reasons for particular characterisations, or explanation of patterns of similarity is often lacking.

Any work using UGC should consider ways in which data quality can impact on interpretations of results. These include properties such as participation inequality, where a small proportion of users contribute a large volume of content (Van Mierlo, 2014), uncertainties in positions or their interpretation (Stvilia & Jörgensen, 2009), factors influencing semantics including ambiguity and automation (Varol, Ferrara, Davis, Menczer, & Flammini, 2017) and underlying behavioural patterns (Sagl, Resch, Hawelka, & Beinat, 2012).

Returning to our starting point and aim—capturing the properties of cities in meaningful ways—we note that many authors have used Lynch’s initial work to explain and justify the choice of UGC. Somewhat surprisingly, perhaps the most complete attempt to replicate the image of the city to date (Filomena, Versteegen, & Manley, 2019) does so based mainly on administrative data (in the form of the road network and building footprints), replicating the potential to perceive through predominantly visual and structural indices. They did however use land use, as determined by OpenStreetMap contributors to capture some place-like properties mainly relating to affordances. Another example is work by Zhang, Zhang, Yu, and Lin (2018) using a collection of images annotated with outdoor objects. They used street network as “a major place for human mobility and activity” to capture and represent one aspect of a place: physical appearance. Others have used the street as a fundamental unit to explain place, for example in Massey’s (1994) seminal work where Kilburn High Street served as an example for the complexity of place, and a more recent study by Capineri (2016) to explore the same street using Flickr photos.

### 3. Data and methods

#### 3.1. Overview

To characterise and compare street level similarity patterns we used two datasets: firstly, a selection of elements from the OpenStreetMap roads layer to characterise paths, and secondly, Flickr metadata capturing the locations, unique user identifiers (UUID), tags and times at which pictures were taken. Before calculating similarities we filtered data to remove biases, and identified relevant salient tags. We calculated similarity between street segments for three dimensions: semantics (which we define as meaning expressed through patterns of tag usage), user behaviour (based on unique user identifiers) and temporal (based on times at which images were taken). We then mapped the most similar street segments and identified a range of characteristic similarity patterns, which we interpret based on the data contributing to the patterns of similarity.

#### 3.2. Modelling paths

To model paths we downloaded the complete OpenStreetMap roads layer provided by Geofabrik<sup>2</sup> within 33 boroughs of Greater London. Geofabrik provides up-to-date packages of OpenStreetMap data for countries and regions. The initial network consisted of a set of ways (ordered sets of nodes) annotated with names, types and references to UK national road classes. We selected only major roads, using the classes primary, trunk and secondary to reduced the density of the overall network and retain important paths. We then removed pseudo-nodes from individual segments with the same name, type and class to form continuous segments where not split by road junctions. Finally, we retained all segments with lengths of more than 200 m, resulting in the street network shown in Fig. 1a. Note that this network is not topologically complete, since short segments were removed. Furthermore, some segments represent individual carriageways of the same street, where these have been digitised as separate segments (e.g., as is the case for expressways with separated lanes). After this process we were left with 3406 unique segments with a median length of 519 m.

#### 3.3. Filtering and assigning attributes to segments

We downloaded an initial Flickr dataset consisting of all georeferenced images available through the Flickr API for the bounding box of Greater London. We then selected only the images found within the polygon representing Greater London, associated with Flickr accuracy [sic] greater than 14. For each image we stored UUID, tags, image coordinates and the timestamp at which a photo was taken. Fig. 1b shows the initial Flickr dataset described here.

Before associating images with street segments, we performed several filtering steps to remove biases typical to UGC and retain salient information. Firstly, we removed images (and users) associated with typical forms of *participation inequality*. We did so by a) removing all users who contributed only a single image (typically not representative tags), b) removing a single very prolific user who contributed some 5% of all images and c) retaining only one image in the case of bulk uploads (i.e., multiple images from one user with identical tags and coordinates). Doing so reduced our initial collection of 5,119,629 images to 2,537,941 images, and the initial 105,021 users to 72,407 users. This filtered dataset, associated with 100 m buffers around street segments, then formed the basis for calculating similarity according to users and time. After extracting only images and users found within the 100 m buffers, we were left with a total of 1,250,205 images and 61,184 users.

Since we wished to calculate and interpret semantic similarity, we not only filtered noise from tags, but also selected semantically relevant

<sup>1</sup> <http://www.sightsmap.com/>

<sup>2</sup> <https://www.geofabrik.de/>

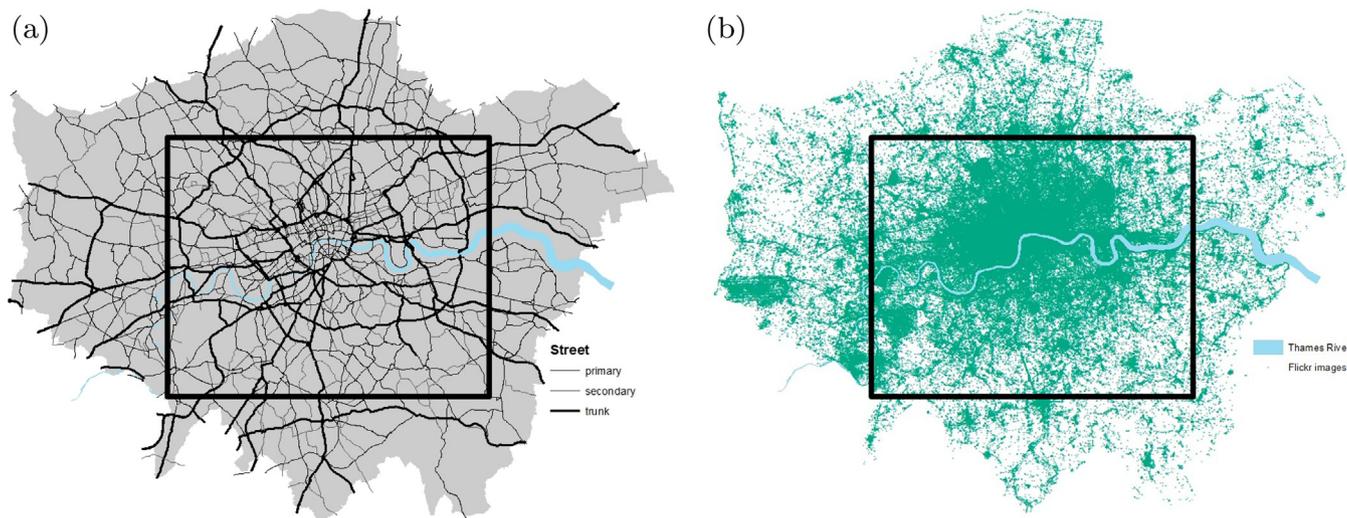


Fig. 1. Study area within 33 boroughs of Greater London.

terms which capture perceived properties of the city. To do so, we removed images with no tags, and tags using non-ASCII characters, duplicate tags in the same list and machine generated tags. We also removed images and tags shared through Instagram links, since we noted that the subjects of such images often had limited relationship to location. Furthermore, many tags used in Instagram relate to memes and filters which are highly ambiguous and biased results (e.g., a popular Instagram filter is called *earlybird*). In previous work [Hollenstein and Purves \(2010\)](#) showed that bias in the use of individual tags could be accounted for by the use of tag profile histograms. These allow us to remove tags with uneven patterns of use (e.g., those used only by a few prolific users). We removed tags with coefficient of variation ( $> 200$ ) from our dataset.

Having filtered tags using these steps, we were left with a total vocabulary of 4744 unique tags and 8,967,337 tags in total describing 1,726,670 images taken by 51,282 users. We then again filtered images to retain only those found within 100 m buffers around street segments. To select the most representative tags from the remaining images, we used Latent Dirichlet Allocation (LDA) to perform topic modelling on a 500 m grid overlaid on our study area ([Bahrehdar & Purves, 2018](#); [Blei, Ng, & Jordan, 2003](#)) as a further filtering step. For each tag, topic modelling outputs the probability of the tag belonging to a particular topic. Based on previous work, we identified 40 topics as an appropriate number to broadly characterise London at a 500 m resolution ([Bahrehdar & Purves, 2018](#)). We then assigned all tags predicting 80% of the cumulative probability per topic to a global list which we retained, thus removing other tags providing limited information about specific locations. Our LDA analysis was carried out using the Machine Learning for Language Toolkit (MALLET) ([McCallum, 2002](#)). We used default hyperparameter values for *alpha* of the reciprocal of the number of topics and *beta* of 0.01 and the 'optimise-interval' option.

Having performed topic modelling, we found a mix of generic terms and proper nouns, in the form of place names, as would be expected from typical tagging behaviour ([Sigurbjörnsson & van Zwol, 2008](#)). Since we did not wish to measure semantic similarity based on place names, but rather properties, we further filtered place names from our tags using fuzzy matching on a set of place names extracted from GeoNames. Finally, we treated the remaining list of tags associated with grid cells as an *allow list* for segments passing through that grid cell. The final filtered dataset created by 36,486 users of 671,207 images described by 1605 unique tags and 4,268,980 tags in total was then used to calculate semantic similarities.

Note that thus two datasets were used in our similarity calculations: one for temporal and user similarities where we did not filter based on

tagging behaviour, and a more strongly filtered dataset for the calculation of semantic similarity.

#### 3.4. Measuring similarities

We calculated **semantic similarity** by comparing tags used to describe segments. Each segment ( $S$ ) is represented as a vector  $V_S = [t_1^s, t_2^s, \dots, t_n^s]$  where each element of the vector  $t_i^s$  corresponds to a tag's frequency in that segment, and  $n$  is the number of unique tags. Since raw counts are biased towards tags which are frequent across London as a whole, we calculated a normalised spatial TF-IDF to increase the weight of tags common in particular segments, but rarer as a whole as follows:

$$tf \cdot idf_{(t_i, s_j)} = ntf_{(t_i, s_j)} \cdot idf_{t_i} \quad (1)$$

where  $ntf$  is the number of times a term ( $t_i$ ) was used associated with a segment ( $s_j$ ) and was normalised based on the number of total terms associated with the segment, and  $idf$  was calculated as follows:

$$idf_{t_i} = 1 + \text{Log}_e \left( \frac{N}{sf_{t_i}} \right) \quad (2)$$

$N$  is the total number of segments and  $sf_{t_i}$  or *segmentfrequency* is the number of segments where term ( $t_i$ ) is found.

Similarity between segment pairs was calculated using cosine similarity for the weighted TF-IDF vectors as the dot product of two vectors:

$$\text{Sim}(v_{s_1}, v_{s_2}) = \cos(\theta) = \frac{v_{s_1} \cdot v_{s_2}}{\|v_{s_1}\| \|v_{s_2}\|} \quad (3)$$

Similarity values of 1 indicate that the semantics of two segments are identical, while values of 0 indicate complete dissimilarity.

To compare how similar two segments are in term of unique users (who have photos associated with segments), we again used cosine similarity. Here, however, we represented each segment as a binary vector, containing either a) all users found in London or b) only those who took images within the segments over two weeks or less. We treated this second group as tourists (c.f. [Girardin et al. \(2008\)](#); [Straumann et al. \(2014\)](#)).

Our fourth similarity dimension was based on the **temporal** distribution of images associated with a segment. We chose to compare segments according to the proportion of visits on different days of the week (c.f. [McKenzie, Janowicz, Gao, and Gong \(2015\)](#)) after experimenting with hours of the day and months of the year. We calculated

temporal similarity as the Euclidean distance between a seven-dimensional vector, where we treated the proportion of images taken on each day of the week as an independent dimension.

### 3.5. Exploring global patterns

Our similarity measures allow each street segment to be compared to all other street segments in London across four dimensions (semantic, users (all and tourists) and temporal). We calculated two global measures to explore ways in which these dimensions captured variation in our data. The first measure captured distinctiveness of segments with respect to each other. As a measure of distinctiveness we calculated the average similarity of each segment to all other segments, then ranked all segments by average similarity. This measure allows extraction of the most and least distinctive segments, and exploration of patterns of similarity in distinctiveness.

A general criticism of work such as ours is that the resulting patterns simply reflect underlying spatial autocorrelation in image distributions. We therefore calculated Global Moran's I for each segment with respect to its similarity to every other segment in all four dimensions. For each segment and dimension we calculated z-scores (indicating the strength and nature of spatial autocorrelation) and p-values (indicating whether the distribution of similarity scores for a particular segment and dimension met our null hypothesis of complete spatial randomness). We considered segments to be spatially autocorrelated where  $p < 0.01$ .

## 4. Results

### 4.1. Global patterns

Fig. 1.a and b show the study area and the segments with which images were associated and the locations of the images analysed. They demonstrate that Flickr images are commonly associated with the network structure of the street network, as represented by our model, but also show concentrations in open spaces (which are not captured).

Table 1 summarises key values related to the distinctiveness of segments. Semantically distinctive segments are moderately correlated with those ranked by users, while all other correlations are low - implying that different segments are distinctive in each dimension. To quantify this effect, we counted the number of shared segments in the 200 most distinctive segments for each dimension. These values demonstrate that even at moderate levels of correlation very different sets of segments are distinctive within the top 200, with a maximum shared set equivalent to only 34% of segments.

Our second global measure explores the spatial autocorrelation of each segment with respect to its similarity to all other segments. Significant values of spatial autocorrelation indicate spatial clustering of similarity values around a particular segment. Using semantics to measure similarity, we found 243 (7%) segments which were spatially autocorrelated with similar segments, with users 2407 (72%) segments, tourists 478 (18%) and temporally only a single segment. These results indicate that users tend to move in particular parts of the city, as do tourists who are confined to a smaller region, resulting in spatially autocorrelated similarity values being relatively common for these dimensions. Fig. 2 shows the locations of these segments, distributed all

**Table 1**

Correlation values  $r^2$  of rank average similarity per segment and number of shared segments in the 200 most distinctive segments for each dimension.

	Semantics	Users	Tourists	Temporal
Semantics	NA	0.373 (68)	0.014 (18)	0.066 (3)
Users	0.373 (68)	NA	0.008 (16)	0.079 (0)
Tourists	0.014 (18)	0.008 (16)	NA	0.028 (18)
Temporal	0.065 (3)	0.079 (0)	0.028 (18)	NA

over London for users and primarily confined to the centre of London north of the river for tourists. For semantics, we observe a different pattern, where spatially autocorrelated tag usage is distributed across London, suggesting that spatial patterns in tag usage are not controlled primarily by image distribution. These initial results confirmed our initial hypothesis that different dimensions allow us to explore segments in different ways, and were not simply correlated with one another. We therefore set out to explore in more detail the properties of four individual segments.

### 4.2. Exploring individual segments

We implemented a Processing tool which allowed us to interactively explore four dimensions simultaneously: semantic, users (all and tourists) and temporal. This tool is available online<sup>3</sup> and it is important to note that the following examples were identified through its use. We describe and interpret the properties of four locations (Table 2), selected because of their contrasting properties and efficacy in illustrating differing aspects of our approach. For example, Tower Bridge was distinctive in terms of tourist behaviour, and showed spatially autocorrelated similarity in semantics, users (as was the case for all four examples) and tourists. Chepstow Road was one of the most distinctive locations temporally. Whitehall was distinctive in terms of semantics and tourists, and again showed spatial autocorrelation in tourist similarity. Crystal Palace Parade was distinctive in terms of semantics and users, but only showed clustering in user behaviour.

The first example is a very well-known London location, Tower Bridge (Fig. 3). We present four maps of correlations between segments and a tag cloud illustrating the segments shared by at least twelve of the thirty most similar segments to Tower Bridge. The most semantically similar segments to Tower Bridge form a sinuous path along the banks of the Thames, linked by many of its bridges. These give the appearance of forming a path through the city sensu Lynch, and exploring the tag cloud reveals that the Thames Path (and Thames River) are indeed tags shared by many of the most similar segments. Many other tags reveal different aspects of this location such as its bridges, boats, tides and the river itself. Some more specific tags, for example, victoriaembankment, theshard, and bankside refer to named locations found along the Thames which were not removed by our toponym filtering. Various image properties, some more likely to be related with water (e.g., reflection and fog) are found, together with a host of photography related terms which could arguably have been filtered (e.g., canon, blackandwhite, nightshot). Nonetheless, our semantic similarity measure both reveals a district, which can also be interpreted as a path, and allows us to interpret it in a meaningful way. Maps of users and tourists show (consistently in four all examples) weaker correlations. Users in general are clustered around Tower Bridge, with a bias to the west, and north of the river, though some users do cross to the south of the Thames. Tourists are found in a smaller region, almost only in Central London either near, or to the north of the river. These maps indicate the effectiveness of the Thames as a barrier to people, with users much less likely to visit seemingly similar regions (as defined through semantics). Temporally, we note that correlations for many segments are high, and see little if any spatial pattern.

Our second example, Chepstow Road (Fig. 4) reveals a spatial pattern of correlated semantic segments, picking out a very small district around Notting Hill. This is the location of the annual Notting Hill Carnival (tagged as nottinghillcarnival), and rather than identifying a district through an affordance (e.g., the banks of the Thames and the Thames Path), here semantic similarity reveals an event. The semantics of the tag cloud reflect this, with shared tags including

<sup>3</sup> Download a zip file [https://www.dropbox.com/s/q2mpr3iczkx1x3i/users\\_cosineSimilarity\\_binary.zip?dl=0](https://www.dropbox.com/s/q2mpr3iczkx1x3i/users_cosineSimilarity_binary.zip?dl=0)



Fig. 2. Locations and Z-scores for segments with statistically significant spatial autocorrelation (Global Moran's I ( $p < .01$ ) for semantics, users and tourists.

**Table 2**  
Selected segments and their global summary values.

Segment	Description		Distinctiveness Rank	Morans I (p-value z-score)
Tower Bridge	An iconic tourist attraction in the centre of London	semantic	1010	(0.00 2.86)
		user	1587	(0.00 7.45)
		tourist	108	(0.00 8.43)
		temporal	1745	(0.61–0.50)
Chepstow Road	Home of the annual Notting Hill Carnival	semantic	1662	(0.34 0.96)
		user	1397	(0.00 6.08)
		tourist	878	(0.76 0.30)
Whitehall	At the political heart of London	temporal	30	(0.99 0.01)
		semantic	239	(0.28 1.08)
		user	987	(0.00 7.21)
Crystal Palace Parade	A suburban south London street	tourist	114	(0.01 2.49)
		temporal	1727	(0.53–0.62)
		semantic	180	(0.33 0.97)
		user	33	(0.00 4.46)
		tourist	1642	(0.76 0.30)
		temporal	1055	(0.75 0.32)

carnival, dancing, parade, party and so on. The pattern of user correlations is more spatially extensive than that for Tower Bridge, revealing that the community visiting this location roams further than that photographing the tourist site of Tower Bridge. Tourists however, appear to share almost no segments in common. Temporal correlations again reveal little variation.

The third example, Whitehall, lies in the heart of London, and is associated with both political and ceremonial events (Fig. 5). Semantically, we can pick out a region around Central London, spanning both sides of the Thames. We note, as was the case for Tower Bridge, a range of tags related to photography and named locations in this region (e.g., *oxfordst*, *stjamespark* and *hydeparkcorner*). Many tags reflect the usage of this part of London, conveying recurring and rare events (e.g., *celebration*, *royalwedding*, *parade*, *protest*) and their participants (e.g., *soldier*, *queen*, *guards*). The users photographing this segment again visit larger areas than those visiting Tower Bridge, with tourists once more focusing on locations north of the river, and the Thames acting as barrier to movement south.

Our final example (Fig. 6), Crystal Palace Parade, reveals a very different pattern to the previous three, all of which allowed us to identify coherent regions associated with semantically similar segments. In this case, these segments are distributed, seemingly randomly, across all of London. However, the tag cloud associated with the most similar segments reveals the reason for this pattern. Other than common tags related to photography, we find here many tags related to transport including bus, types of bus (e.g., *scania*, *plaxton*, *rou-temaster*, *mercedes*, *volvo*) and providers of public transport (e.g., *arriva*, *londontransport*, *stagecoach*, *abellio*). Semantic similarity with this location is thus defined by photographs of a particular

type, taken by a specialist group interested in public transport. Users present at this location spread not only over south London, but into north London as well, demonstrating an asymmetry in the barrier effect of the Thames apparently limiting movement from north to south more than south to north. Since semantic similarity and user similarity have very different patterns, our method implicitly shows that different photographers are interested in the same subject matter. Tourists taking pictures at this location appear to be rare, and thus have a very limited local spatial spread. Since this pattern, and subject matter, was unexpected, we performed a Google Image search using “Crystal Palace Parade” as a search term. Seven of the top 20 images indeed contained buses, suggesting that our approach had indeed identified a semantically meaningful pattern.

Having explored these individual examples, the obvious question which arises is, how can we interpret these results more generally, and can the results be linked to the ideas posed by Lynch? With respect to the former question, we note that by using three distinct dimensions (semantics, users and time) different patterns are revealed. The patterns associated with users form regions or districts sensu Lynch clustered around the query segment in all cases, though the forms of these regions are not always symmetrical. This effect mirrors the large number of segments whose patterns of similarity we found to be spatially autocorrelated in the user dimension. Thus, for Tower Bridge we note a general tendency to locations north of the river, revealing how the Thames works as a barrier, or in Lynch's terms an edge. However, at Crystal Palace, south of the river, this barrier is less influential, revealing a different pattern of user behaviour—users here appear less influenced by the Thames as a barrier or edge than those to the north. When selecting out tourists alone, based on their length of stay in London, we find meaningful signatures (which largely replicate the pattern of users in general) only at very popular sites (e.g., Tower Bridge and Whitehall). Our semantic signatures are interesting in a number of different ways. Firstly, they reveal not only where similar aspects of a scene were annotated, but also what was of interest. These include named locations (landmarks sensu Lynch) as well as objects commonly found in scenes and properties of scenes. Each example has quite different semantic properties, and the form of the districts associated with similar semantics vary from the linear path through London generated by the Thames and the Thames Path for Tower Bridge, through Central London as a whole associated with Whitehall, to the very small region related to the Notting Hill Carnival for Chepstow Road, and finally the dispersed locations associated with public transport for Crystal Palace, where no meaningful district emerges. As we saw globally, semantic similarity is less often spatially autocorrelated, and this is mirrored here. Temporally, our method struggles to identify similar regions since the overall distribution of Flickr images shows limited temporal variation.

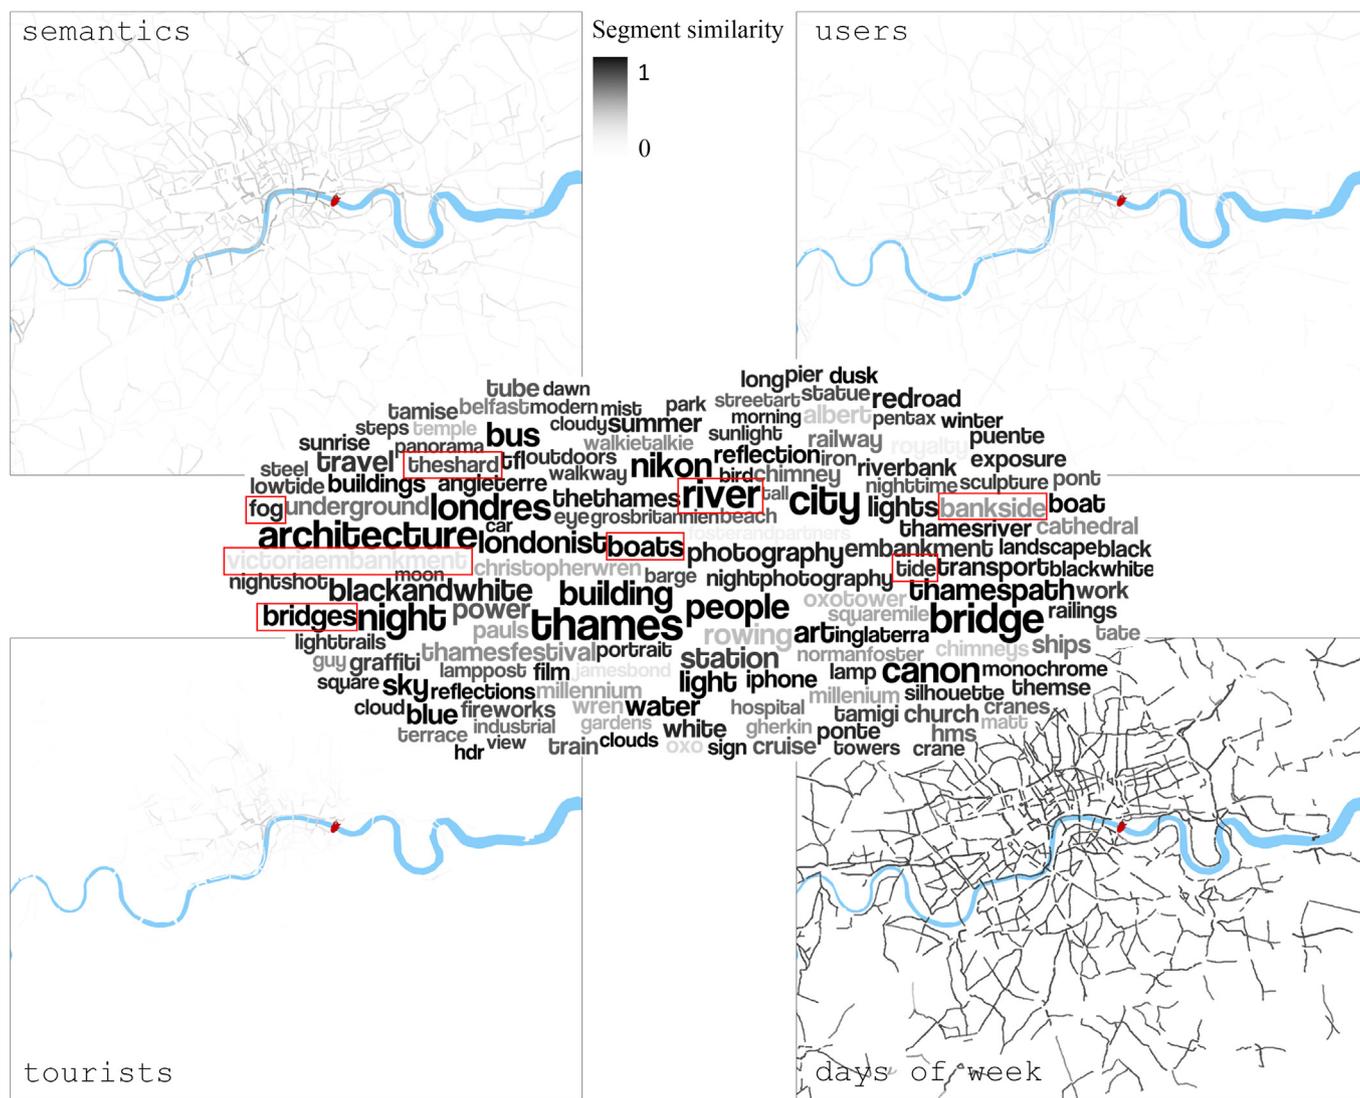


Fig. 3. Signature similarities for Tower Bridge: Each of four maps represents the similarity between the queried street in red and all other streets in London. Darker segments are the more similar. The word cloud presents the shared semantics tag of the 30 most similar segments. Larger tags are higher ranked, and darker tags are shared by more of the top 30 segments. Tags highlighted in red are discussed in the text. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

### 5. Discussion

We set out to develop a tool which would allow us to capture information about how a city was perceived through the properties of an element identified by Lynch as central to our understanding—paths through the city. By associating UGC, in the form of tags, timestamps and UUIDs with street segments, we were able to interactively explore similarity between segments across three, contrasting, dimensions. Exploring global properties in terms of distinctiveness and spatial autocorrelation allowed us to demonstrate that the different dimensions, at least globally, revealed different patterns and were worthy of further exploration.

In the following, we firstly discuss key influences on, and limitations of, our approach, and discuss it with respect to previous work, before setting out our contribution in the broader context of practice.

The first, and most important influence concerns the filtering of our data. We chose to filter based both on behaviour (e.g., taking account of participation inequality, bulk uploads and so on), semantic biases (primarily seeking to retain only tags used by a broad group of users) and identify semantically distinctive terms (through topic modelling). These choices mean that we explore the temporal and user dimensions

with different data sets to the semantic one; however, we argue that knowingly making these choices is a valid approach. Although filtering is often left implicit, or only briefly discussed, in our case these choices reduced the original dataset five-fold. We think the implications and importance of filtering have been neglected in the gold-rush mentality of analysis of UGC, and believe that the attention now being paid to bias in data in artificial intelligence tasks (Zou & Schiebinger, 2018) is equally important here (Shelton et al., 2014).

An important limitation is our approach to filtering using Geonames to remove place names from lists of tags. Not all place names are captured in Geonames, especially at finer granularities (e.g. *bankside* and *theshard*), and some compound tags also remain (e.g. *victoriaembankment* and *thamesriver* in Fig. 3).

In linking tags to segments we chose a buffer width of 100 m; changing this width would also reduce or increase the number of image locations associated with segments. Increasing buffer size would however reduce distinctiveness of tags, since they would be associated with multiple segments, while smaller buffers would lead to a very limited set of tags for less well-covered regions outside of Central London. We explored the overlap of tags between very similar segments, and found that, for example, for the ten most temporally similar segments to



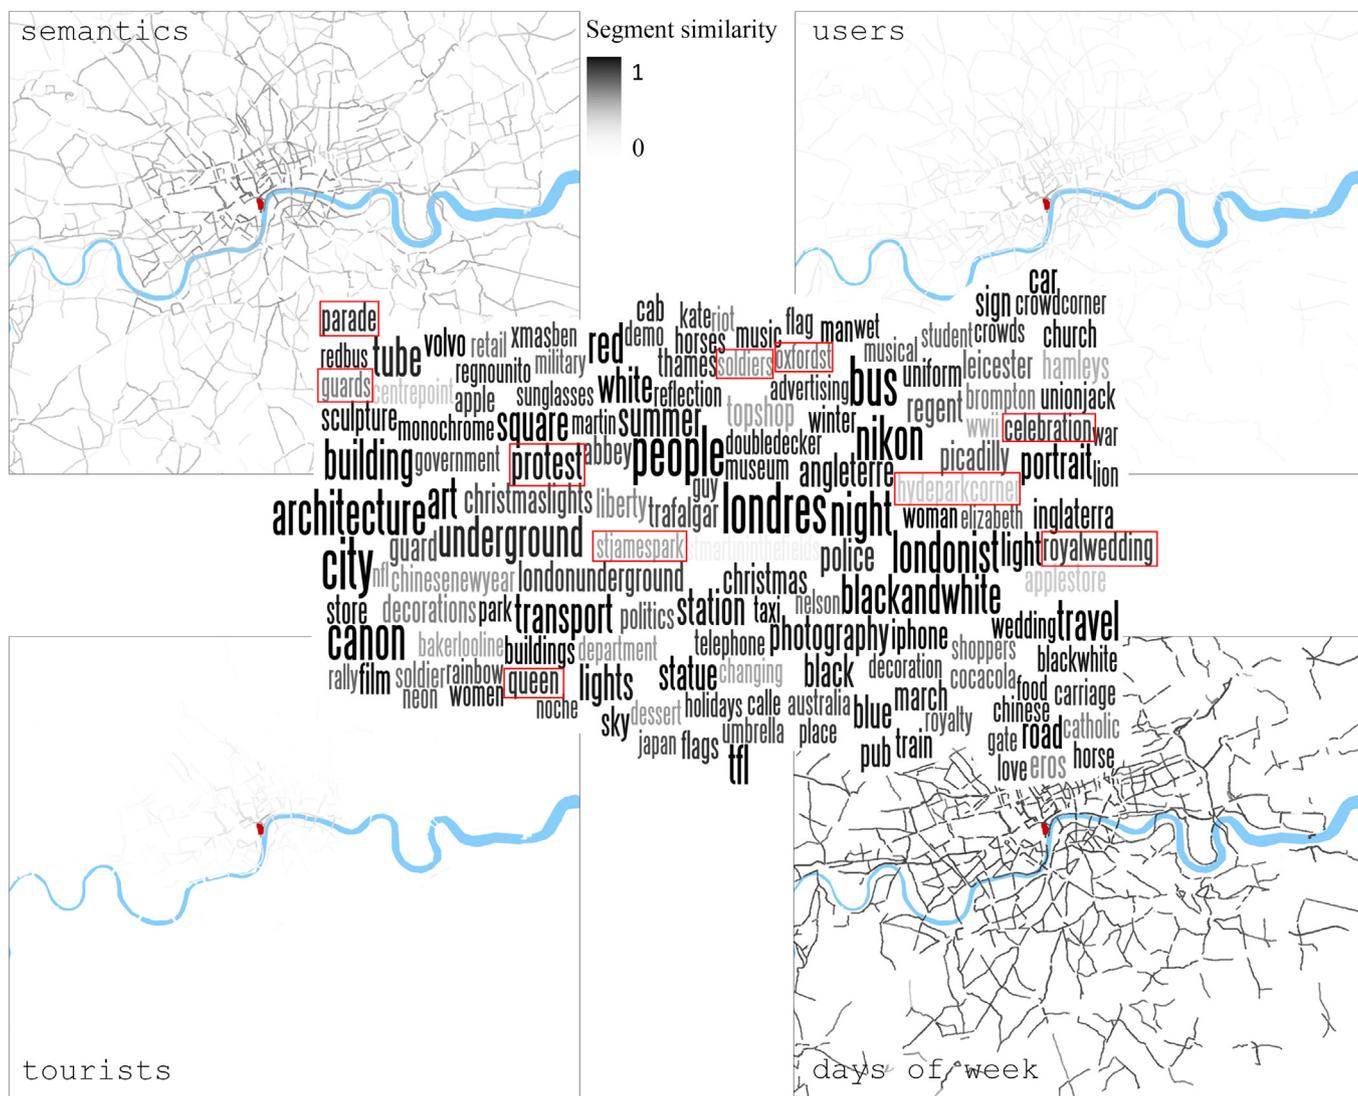


Fig. 5. Signature similarities for Whitehall: Each of four maps represents the similarity between the queried street in red and all other streets in London. Darker segments are the more similar. The word cloud presents the shared semantics tag of the 30 most similar segments. Larger tags are higher ranked, and darker tags are shared by more of the top 30 segments. Tags highlighted in red are discussed in the text. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

paths, we can not only find similar regions, but describe their properties and link these to behaviour in the city itself. However, it is important to note that the missing parts of the city, where we find no data, are potentially just as important in understanding how the city is perceived by its inhabitants, and our approach, and others focussing on passive crowdsourced data cannot address this gap. Active approaches such as those proposed by the mappiness app (Seresinhe, Preis, MacKerron, & Moat, 2019) may go some way to filling this hole, but the importance of such data gaps cannot be overstated (Graham, Hogan, Straumann, & Medhat, 2014). Nonetheless, our approach starts to suggest how Lynch's ideas can be empirically implemented at scale.

According to Lynch ([p. 8]1960), a workable image of a city requires three important elements such as identity (in the sense of identification of urban elements), structure (indicating spatial or pattern relation among urban elements, for example, in a street network), and meaning (either practical or emotional meaning for an observer). By capturing multiple dimensions of similarity, and linking these to paths through the city, analysis not only of space, but place is enabled, and in doing so important relationships between locations are revealed. Our approach allows, in principle, exploration across time steps, and thus is temporally dynamic, and synthesises heterogeneous data. The tool is

easy (and we think fun!) to use, and interactivity enhances exploration. We note that our dimensions could also be combined, exploring for example semantic similarity at particular times, or for particular user groups, though doing so would require that the same filtering approach was taken with all data.

### 6. Conclusion and future work

Starting with Dick Whittington's confusion when confronted with a London very different to the stories he had heard, we set out to model the characteristics and thus similarity between streets in London using user generated content. Streets are a natural unit, since they capture the paths described by Lynch, and our results demonstrate how they allow us to explore perception in terms of not only paths through the city, but through the emergence of districts, landmarks and even edges. These elements emerge because we explore different dimensions capturing semantic similarity, user behaviour and temporal patterns. Street segments are a more natural way of organising data, and reduce the issues caused by aggregating across administrative boundaries or arbitrarily imposed tessellations such as grids. We demonstrate that the data found in London are sufficiently rich, despite numerous filtering steps, to



**Fig. 6.** Signature similarities for Crystal Palace Parade: Each of four maps represents the similarity between the queried street in red and all other streets in London. Darker segments are the more similar. The word cloud presents the shared semantics tag of the 30 most similar segments. Larger tags are higher ranked, and darker tags are shared by more of the top 30 segments. Tags highlighted in red are discussed in the text. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

reveal interesting and meaningful patterns, though interpretation of these requires us to both drill down into the data and use external knowledge. Semantics and user behaviour contain more variation in Flickr data than temporal patterns, showing how the the choice of dimension is important in exploring a city.

Exploring global patterns in each dimension was a useful starting point to understand the general properties of our data—for example, user behaviour was spatially autocorrelated for a much larger proportion of segments than semantics, demonstrating that the way in which parts of London are described is not only influenced by who describes London. Local patterns were explored for four contrasting locations, identified through exploration using an interactive tool. These locations demonstrate that not only touristically important central locations such as Tower Bridge and Whitehall are captured in UGC, but also events (Notting Hill Carnival at Chepstow Road) and the association of locations with particular interest groups in less touristic areas (as shown for Crystal Palace Parade).

We suggest that future work aiming to use UGC in planning or applications such as location based services consider how such data can be effectively integrated, while not forgetting the implications of data bias and gaps. In particular, we propose integrating data sources which

capture missing or underrepresented dimensions here. For instance, Twitter has been shown to capture temporal patterns in London (Lansley & Longley, 2016).

**Funding**

Notes on contributor(s).

**Declaration of Competing Interest**

None.

**Acknowledgements**

This work was supported by the Schweizerischer Nationalfonds zur Förderung der Wissenschaftlichen Forschung [200021\_149823] and the New Zealand Building Better Homes, Towns and Cities National Science Challenge.

## References

- Adams, Benjamin, & McKenzie, Grant (2013). Inferring thematic places from spatially referenced natural language descriptions. *Crowdsourcing Geographic Knowledge* (pp. 201–221). .
- Bahrehdar, A. R., & Purves, R. S. (2018). Description and characterization of place properties using topic modeling on georeferenced tags. *Geo-spatial Information Science*, 21(3), 173–184.
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3, 993–1022.
- Boy, J. D., & Uitermark, J. (2017). Reassembling the city through Instagram. *Transactions of the Institute of British Geographers*, 42(4), 612–624.
- Capineri, C. (2016). Kilburn high road revisited. *Urban Planning*, 1(2), 128–140.
- Carmona, M., Heath, T., Oc, T., & Tiesdell, S. (2012). *Public places-urban spaces*. Routledge.
- Crandall, David J., Backstrom, Lars, Huttenlocher, Daniel, & Kleinberg, Jon (2009). Mapping the World's Photos. *Proceedings of the 18th International Conference on World Wide Web, WWW '09, New York, NY, USA* (pp. 761–770). .
- Derungs, C., & Purves, R. S. (2016). Characterising landscape variation through spatial folksonomies. *Applied Geography*, 75, 60–70.
- Dunkel, A. (2015). Visualizing the perceived environment using crowdsourced photo geodata. *Landscape and Urban Planning*, 142, 173–186.
- Elwood, Sarah, Michael F Goodchild, and Daniel Sui. 2013. "Prospects for VGI research and the emerging fourth paradigm." In *Crowdsourcing Geographic Knowledge*, 361–375. Springer.
- Feick, R., & Robertson, C. (2015). A multi-scale approach to exploring urban places in geotagged photographs Special Issue on Volunteered Geographic Information. *Computers, Environment and Urban Systems*. 53. *Computers, Environment and Urban Systems* (pp. 96–109).
- Filomena, G., Versteegen, J. A., & Manley, E. (2019). A computational approach to the image of the City. *Cities*, 89, 14–25.
- Fischer, E. (2013). *The geotaggers world atlas*.
- Gao, S., Janowicz, K., Montello, D. R., Hu, Y., Yang, J.-A., McKenzie, G., ... Yan, B. (2017). A data-synthesis-driven method for detecting and extracting vague cognitive regions. *International Journal of Geographical Information Science*, 31(6), 1245–1271.
- Girardin, F., Calabrese, F., Fiore, F. D., Ratti, C., & Blat, J. (2008). Digital footprinting: Uncovering tourists with user-generated content. *IEEE Pervasive Computing*, 7(4), 36–43.
- Goodchild, M. F. (2011). Formalizing place in geographic information systems. *Communities, neighborhoods, and health* (pp. 21–33). Springer.
- Graham, M., Hogan, B., Straumann, R. K., & Medhat, A. (2014). Uneven geographies of user-generated information: Patterns of increasing informational poverty. *Annals of the Association of American Geographers*, 104(4), 746–764.
- Hahmann, S., Purves, R., & Burghardt, D. (2014). Twitter location (sometimes) matters: Exploring the relationship between georeferenced tweet content and nearby feature classes. *Journal of Spatial Information Science*, 2014(9), 1–36.
- Hauthal, E., & Burghardt, D. (2013). Detection, analysis and visualisation of georeferenced emotions. *Proceedings of the 26th International Cartographic Conference*.
- Hobel, Heidelinde, Paolo Fogliaroni, and Andrew U Frank. 2016. "Deriving the geographic footprint of cognitive regions." In *Geospatial Data in a Changing World*, 67–84. Springer.
- Hollenstein, L., & Purves, R. (2010). Exploring place through user-generated content: Using Flickr tags to describe city cores. *Journal of Spatial Information Science*, 2010(1), 21–48.
- Hospers, G.-J. (2010). Lynch's the image of the City after 50 years: City marketing lessons from an urban planning classic. *European Planning Studies*, 18(12), 2073–2081.
- Hu, Y., Gao, S., Janowicz, K., Yu, B., Li, W., & Prasad, S. (2015). Extracting and understanding urban areas of interest using geotagged photos. *Computers, Environment and Urban Systems*, 54, 240–254.
- Huang, H. (2016). Context-aware location recommendation using geotagged photos in social media. *ISPRS International Journal of Geo-Information*, 5(11), 195.
- Jenkins, A., Croitoru, A., Crooks, A. T., & Stefanidis, A. (2016). Crowdsourcing a collective sense of place. *PLoS One*, 11(4), Article e0152932.
- Keim, D. A., Mansmann, F., Schneidewind, J., Thomas, J., & Ziegler, H. (2008). Visual analytics: Scope and challenges. *Visual Data Mining* (pp. 76–90). Springer.
- Kennedy, L., Naaman, M., Ahern, S., Nair, R., & Rattenbury, T. (2007). How flickr helps us make sense of the world: context and content in community-contributed media collections. *Proceedings of the 15th ACM international conference on Multimedia* (pp. 631–640). ACM.
- Lansley, G., & Longley, P. A. (2016). The geography of twitter topics in London. *Computers, Environment and Urban Systems*, 58, 85–96.
- Lynch, K. (1960). *The image of the city*. Vol. 11. MIT press.
- Massey, D. (1994). A global sense of place. *Space, Place, and Gender* (pp. 146–156). University of Minnesota: Press.
- McCallum, Andrew Kachites (2002). *Mallet: A machine learning for language toolkit*.
- McKenzie, G., Janowicz, K., Gao, S., & Gong, L. (2015). How where is when? On the regional variability and resolution of geosocial temporal signatures for points of interest. *Computers, Environment and Urban Systems*, 54, 336–346.
- Miller, H. J. (2010). The data avalanche is here. Shouldnt we be digging? *Journal of Regional Science*, 50(1), 181–201.
- Miller, R. B., & Small, C. (1999). Digital cities. I. Integrating data and information resources, towards digital Earth. *Proceedings of the International Symposium on Digital Earth* (pp. 217–222). Beijing: Science Press.
- Montello, D. R. (2003). Regions in geography: Process and content. *Foundations of Geographic Information Science* (pp. 173–189). .
- Quercia, Daniele, Schifanella, Rossano, & Aiello, Luca Maria (2014). The Shortest path to happiness: recommending beautiful, quiet, and happy routes in the city. *Proceedings of the 25th ACM Conference on Hypertext and Social Media, HT '14, New York, NY, USA* (pp. 116–125). .
- Quercia, D., Schifanella, R., Aiello, L. M., & McLean, K. (2015). Smelly maps: the digital life of urban smellscapes. *Proceedings of the Ninth International AAAI Conference on Web and Social Media*, 19, 327–336. <https://www.aaai.org/ocs/index.php/ICWSM/ICWSM15/paper/viewFile/10572/10516>.
- Rattenbury, T., & Naaman, M. (2009). Methods for extracting place semantics from Flickr tags. *ACM Transactions on the Web*, 3(1), 1–30.
- Rorissa, A. (2008). User-generated descriptions of individual images versus labels of groups of images: A comparison using basic level theory. *Information Processing & Management*, 44(5), 1741–1753.
- Sagl, G., Resch, B., Hawelka, B., & Beinat, E. (2012). From social sensor data to collective human behaviour patterns: Analysing and visualising spatio-temporal dynamics in urban environments. *Proceedings of the GI-Forum* (pp. 54–63). Berlin: Herbert Wichmann Verlag.
- von Schinfeld, K. C., & Bertolini, L. (2017). Urban streets: Epitomes of planning challenges and opportunities at the interface of public space and mobility. *Cities*, 68, 48–55.
- Seresinhe, C. I., Preis, T., MacKerron, G., & Moat, H. S. (2019). Happiness is greater in more scenic locations. *Scientific Reports*, 9(1), 4498.
- Shelton, T., Poorthuis, A., Graham, M., & Zook, M. (2014). Mapping the data shadows of hurricane Sandy: Uncovering the sociospatial dimensions of big data. *Geoforum*, 52, 167–179.
- Sigurbjörnsson, B., & van Zwol, R. (2008). Flickr tag recommendation based on collective knowledge. *Proceedings of the 17th International Conference on World Wide Web* (pp. 327–336). New York, NY, USA: WWW '08. ACM <https://doi.org/10.1145/1367497.1367542>.
- Straumann, R. K., Cöltekin, A., & Andrienko, G. (2014). Towards (re) constructing narratives from georeferenced photographs through visual analytics. *The Cartographic Journal*, 51(2), 152–165.
- Stvilia, B., & Jörgensen, C. (2009). User-generated collection-level metadata in an online photo-sharing system. *Library & Information Science Research*, 31(1), 54–65.
- Tversky, B., & Hemenway, K. (1983). Categories of environmental scenes. *Cognitive Psychology*, 15(1), 121–149.
- Van Mierlo, T. (2014). The 1% rule in four digital health social networks: An observational study. *Journal of Medical Internet Research*, 16(2).
- Varol, O., Ferrara, E., Davis, C., Menczer, F., & Flammini, A. (2017). *Online Human-Bot Interactions: Detection, Estimation, and Characterization*. 280–289.
- van Weerdenburg, D., Scheider, S., Adams, B., Spierings, B., & van der Zee, E. (2019). Where to go and what to do: Extracting leisure activity potentials from web data on urban space. *Computers, Environment and Urban Systems*, 73, 143–156.
- Zhang, F., Zhang, D., Yu, L., & Lin, H. (2018). Representing place locales using scene elements. *Computers, Environment and Urban Systems*, 71, 153–164.
- Zheng, X., Han, J., & Sun, A. (2018). A survey of location prediction on twitter. *IEEE Transactions on Knowledge and Data Engineering*, 30(9), 1652–1671.
- Zielstra, D., & Hochmair, H. H. (2013). Positional accuracy analysis of Flickr and Panoramio images for selected world regions. *Journal of Spatial Science*, 58(2), 251–273.
- Zou, J., & Schiebinger, L. (2018). AI can be sexist and racist-it's time to make it fair. *Nature*, 559(7714), 324.