

Automatic Extraction of Urban Outdoor Perception from Geolocated Free-Texts

Frances A. Santos · Thiago H. Silva ·
Antonio A. F. Loureiro · Leandro A.
Villas

Received: date / Accepted: date

Abstract The automatic extraction of urban perception shared by people on location-based social networks (LBSNs) is an important multidisciplinary research goal. One of the reasons is because it facilitates the understanding of the intrinsic characteristics of urban areas in a scalable way, helping to leverage new services. However, content shared on LBSNs is diverse, encompassing several topics, such as politics, sports, culture, religion, and urban perceptions, making the task of content extraction regarding a particular topic very challenging. Considering free-text messages shared on LBSNs, we propose an automatic and generic approach to extract people's perceptions. For that, our approach explores opinions that are spatial-temporal and semantically similar. We exemplify our approach in the context of urban outdoor areas in Chicago, New York City and London. Studying those areas, we found evidence that LBSN data brings valuable information about urban regions. To analyze and validate our outcomes, we conducted a temporal analysis to measure the results' robustness over time. We show that our approach can be helpful to better understand urban areas considering different perspectives. We also conducted a comparative analysis based on a public dataset, which contains volunteers' perceptions regarding urban areas expressed in a controlled experiment. We observe that both results yield a very similar level of agreement.

Frances A. Santos
University of Campinas, Brazil
E-mail: frances.santos@ic.unicamp.br

Thiago H. Silva
Universidade Tecnológica Federal do Parana, Brazil
Current address: University of Toronto, Canada
E-mail: th.silva@utoronto.ca

Antonio A. F. Loureiro
Federal University of Minas Gerais, Brazil

Leandro A. Villas
University of Campinas, Brazil

Keywords Perception extraction · Social media · Text mining · Urban outdoor areas · Dictionary creation.

1 Introduction

Cities are not just a place composed of buildings, streets and people living together, but also a place where people have different experiences. It is known that the visual quality of urban areas and other characteristics, such as crime rate, and noise level, can evoke different perceptions about the urban area (Nasar, 1990). Places with tourist attractions, for example, can be appreciated by tourists but avoided by residents performing their daily routines. In this direction, studies have shown that urban perception may affect people’s behavior (Ross and Mirowsky, 2001; Keizer et al., 2008). In Section 2, we introduce the concept of urban perception and also how it might be important to help us understand the latent aspects of urban areas.

Capturing urban perception is not an easy task. One of the challenges is to obtain appropriate data. Field surveys and sensory walks are traditional tools to collect the opinions of people about places (Nasar, 1990; Henshaw, 2013; Quercia et al., 2015). Nevertheless, those strategies could be time-consuming because they usually demand a high amount of time to interview and observe volunteers and collect a considerable amount of perception samples. Thus, those strategies make it difficult to perform this type of analysis for a high number of urban areas. We explore location-based social network (LBSN) data to tackle this challenge, like other recent studies discussed in Section 2.

The number of people that use LBSNs to share their experiences, knowledge, criticism, and opinions on the most diverse topics is expressive (Smith and Anderson, 2018; Shaban, 2019). Thus, LBSNs become an interesting choice to capture urban perception. The content is rich, enabling us to find general comments on almost everything. However, while LBSNs can help in traditional approaches’ scalability issues, their use for urban perception extraction is not simple. Obtaining useful urban perceptions shared by individuals on LBSNs is challenging.

Researchers have shown evidence that LBSNs offer relevant urban perceptions (Quercia et al., 2015; Aiello et al., 2016). Nevertheless, these previous efforts consider LBSN data to extract perceptions about some specific aspects, such as the smell of the environment (Quercia et al., 2015) and urban sounds (Aiello et al., 2016). Moreover, previous studies combine traditional approaches, such as sensory walks, with LBSN data to extract the urban perceptions and, thus, lack flexibility and face scalability issues. Our work tackles those issues by offering an approach to extract urban perceptions from free-texts, without demanding traditional methods, such as sensory walks. Besides, it could be adapted to work on a different subject while still using only public LBSN content. We demonstrate our approach by considering real data from

public reviews of places made on Google Places¹, and Foursquare², as well as public tweets from Twitter. All these data refer to Chicago and New York City, USA, and London, UK. With that, we show our approach’s potential to understand the intrinsic characteristics of areas according to the people’s perception.

The main contributions of this study are:

- An automatic strategy to capture some of the most critical perceptions related to urban outdoor areas from geolocated text. An essential step of our approach is to construct a dictionary, namely UOP-dictionary, which organizes the main descriptive words (i.e., adjectives) used by people to qualify their experiences in urban outdoor areas of cities (Section 3).
- An unsupervised clustering algorithm to identify content shared by individuals in free-text messages that are spatial-temporal and semantically similar. This procedure is fundamental to uncover collective perceptions of groups of people who have visited the same place in the same period (Section 4).
- A set of experiments to evaluate the extracted perceptions of different urban outdoor areas. We have used a Twitter dataset to demonstrate the potential of our approach to uncover the urban perceptions of outdoor spaces that emerge from LBSNs. Studying several scenarios based on LBSN data, we found valuable insights indicating the potential to better understand urban areas in many aspects. We contrasted our results with Place Pulse 2.0, a controlled experiment that contains volunteers’ perceptions expressed in different urban outdoor areas (Dubey et al., 2016). We observe that our approach yields results very similar to those indicated by Place Pulse (Section 5).

2 Contextualization and Related Work

Human perception is a complex process that relies on several factors, such as culture and age. There are many definitions of perception; for instance, Robbins (2003) defines it as the process by which individuals organize and interpret their sensory impressions to give meaning to their environments. Pareek (1988) explains the same concept as the process of receiving, selecting, organizing, checking, and reacting to sensory stimuli or data.

While there are several definitions, we agree that a key aspect of perception is the process of interpreting sensory stimulus, coming from the five senses: sight, taste, smell, touch, and sound. Thus, the sensed information is sent to our brain. Then the perception is how we interpret the sensed information (or sensations) to make sense of everything in our surroundings. In this way, the perception is related to sensation but is much more than that.

¹ <https://cloud.google.com/maps-platform/places>. Last accessed on 12 September 2020.

² <https://pt.foursquare.com/trip-tips>. Last accessed on 12 September 2020.

Each individual can perceive the same place, object or situation differently based on relevant aspects to him/her. The main reason for this is that people interpret things by considering their past experiences and the values they consider appropriate. The perceptions of a group can also influence one's perception. Thus, perceptions can differ from person to person.

We can obtain individuals' perceptions of virtually anything. However, in this study, aligned with others, we focus on the individuals' perception of urban environments, called *urban perception*. More specifically, we are interested in perceptions regarding urban outdoor areas. Those types of areas, such as parks, streets, and plazas, may offer people the opportunity of having diverse experiences, and, for this reason, could trigger different perceptions among their visitors.

A traditional approach to capture urban perceptions is sensory walks, where volunteers answer surveys that contribute to building the common sense about inherent characteristics of cities, which is a common approach to collect features regarding urban areas (Keizer et al., 2008; Henshaw, 2013). This approach provides fine-grained data about urban spaces' perceptions, but it does not scale easily (Quercia et al., 2015; Dubey et al., 2016). To overcome this problem, a possible strategy is to extract urban characteristics shared by people in online sources, such as crowdsourcing systems and location-based social networks (LBSNs), which are useful for understanding and tackling different problems.

Using LSBN data, we can study a variety of urban outdoor perceptions, ranging from perceived safety and traffic conditions to aesthetic characteristics and the cost of living in urban areas. Some studies concentrate their efforts to understand the factors present in the environment responsible for triggering such perceptions.

Towards that direction, Quercia et al. (2014) collected people's perceptions about photos taken in the streets of London. They found a correlation among colors, texture and visual patches of those photos with beautiful, happiness, and quiet perceptions. In the same direction, Naik et al. (2014) proposed an approach for predicting the perceived safety of cities exploring images from different sites. Similarly, Porzi et al. (2015) and De Nadai et al. (2016) demonstrated that computer vision techniques could also identify human perception and predict judgments of safety from images of urban scenes. Dubey et al. (2016) explored a global dataset of urban images to rank street-level images of city aesthetics by also using computer vision techniques. Their results showed that urban perception data on a worldwide scale could be extracted from online photos.

Quercia et al. (2015) explored sensory walks to collect the citizens' perception concerning the smell of the environment. This offline process has enabled the creation of a dictionary with urban-smell related words, which is used to discover messages related to odor perceptions into social media data. Quercia et al. (2016) presented a follow-up of this study, where the authors investigated, among other things, the relationship between the predominant color of the image (visual perception) and the smell associated with the image (olfac-

tory perception). Also, considering the citizens' perception regarding urban smell, Hsu et al. (2019) developed an online crowdsourcing system to enable Pittsburgh's community to report odors and uncover where those odors are frequently concentrated.

Aiello et al. (2016) explored the influence that urban sounds have on the way people perceive places. The authors created a dictionary with sound-related words to discover the urban sounds responsible for triggering the people's perceptions and emotions, which was used to mine LBSN data related to urban sounds. In this way, London and Barcelona's streets could be mapped with one of the six considered categories (transport, mechanical, nature, human, music, indoor) of the dictionary created.

In the literature, different modeling strategies have been used to extract information about a given subject. Several proposals have characterized the urban spaces according to socio-cultural activities (Steiger et al., 2016), functional zones (Yuan et al., 2015; Jenkins et al., 2016; Yao et al., 2017), semantics and sentiments (or emotions) (Hu et al., 2019), or points of interest (Jiang et al., 2016).

Recent studies have made efforts to automatically uncover and predict the perception reflected by urban areas, using, to this end, pictures of outdoor places collected from online sources. Santani et al. (2018) collected, using mobile crowdsourcing, a total of 7,000 geotagged images from three Mexican cities, which were used to gather opinions from 144,000 "workers" of the Amazon Mechanical Turk. With this, the authors proposed a deep learning framework to learn and characterize the urban perception based on visual features present in the observed images and the perceptions related to them. Similarly, Redi et al. (2018) collected geolocated Flickr pictures, and the tags attached to them, posted by individuals in Greater London. Next, the authors proposed a methodology to map London's neighborhoods' ambiance, based on a psychological taxonomy.

As some of the related studies, we also take advantage of LBSN data to facilitate the process of understanding people's perception of urban areas. We acknowledge that using only social media as a data source can be problematic in some cases, such as in smaller cities and less tourism-related places, where the extraction of relevant perception, if they are identified, can fall only within some specific topics, perhaps less informative. However, despite considering only social media data in our work, due to several reasons (e.g., easy access, public availability, large-scale data for larger cities, and rich content), the proposed approach does not depend on any exclusive social media characteristics to properly work, as other literature studies that use, for instance, hashtags, emoticons, and tags. In this sense, we consider that other kinds of data sources, such as IoT devices, smart vehicles, and government data, could be used by our approach. They can have natural language texts, geospatial, and temporal information, potentially contributing to deal with bias present in social media data.

Moreover, our work significantly differs from previous studies in several aspects. For instance, the format in which urban perceptions can be shared on

LBSN is diverse, such as via text, image, video and audio. We focus our work only on perceptions shared via text message. The proposed approach is also one of the very few ones that explore free-texts shared by users on LBSNs to extract several perceptions related to urban outdoor areas. Our approach does not require time-consuming field surveys or manual steps to extract urban outdoor perceptions.

Finally, our approach also differs from social media specialized in rating places, such as Yelp, TripAdvisor, Foursquare and Google Places. These systems enable their users to share reviews about places, containing rich information regarding our area of study. However, data on such systems are usually more available to indoor places, such as restaurants, and points of interest, such as museums, rather than the more general urban areas, such as avenues, streets and squares. Thus, opinions unlinked to an establishment, i.e., perceptions shared while transiting in the city, may not be properly captured. Besides, they are heavily dependent on users' engagement who must actively participate in these systems; thus, for some unpopular places, there might be only a few, if any, data. Exploring only these sources might offer a partial view of the whole urban perception spectrum. In this way, our approach can be helpful to better understand urban areas considering different perspectives, in an automatic and generic way.

3 Dictionary Creation

The use of tweets to extract urban perception is difficult given their content diversity, which is not restricted to the subject under investigation. To overcome this problem is essential first to explore a less noisy source and, thus, learn properly the vocabulary used by people regarding outdoor areas.

In this sense, we selected outdoor urban areas in Chicago and New York City (NYC), United States, and London, United Kingdom, to collect public reviews shared before February 2017 on Google Places and Foursquare (Tips). Such systems enable us to specify the place categories^{3,4}, which allow selecting just reviews about outdoor places. Additionally, we filtered out reviews not written in English. We call this dataset of *Places Review*. Table 1 summarizes the number of collected reviews. As we can see, due to the restriction of Google Places API, the documents from Foursquare represents most of the *Places Review* dataset, which has a total of 39,348 reviews. Using this dataset, we want to uncover frequently used words to qualify experiences in urban outdoor areas. This step is essential to build our urban outdoor perception dictionary, namely the UOP-dictionary.

More formally, the *Places Review* dataset can be defined as follows:

³ Foursquare Venue Category Hierarchy: <https://developer.foursquare.com/docs/build-with-foursquare/categories/>. Last accessed on 12 September 2020.

⁴ Google Places Category: https://developers.google.com/places/web-service/supported_types. Last accessed on 12 September 2020.

Table 1 Reviews in the *Places Review* dataset.

	Chicago	London	NYC
Foursquare Tips	5,085	7,261	24,921
Google Places	666	662	753

Definition 1 A collection of documents \mathcal{D}_R , where each $doc \in \mathcal{D}_R$ is a document determined by a tuple $doc = (id, s, \tau)$, where id is a unique identifier, s is a list of sentences that comprise all content written by the user, and $\tau \in \mathbb{R}$ is a timestamp.

In this work, the term sentence is used to refer to the preprocessed free-text, where numbers, Uniform Resource Locators (URLs), special characters, punctuation, and stop words were removed, and frequent English contractions, e.g., *aren't* and *I 'm*, were replaced by the expanded form, e.g., *are not* and *I am*. Next, we perform the stemming process to obtain the words in the common root form⁵, resulting in a vector of single words' stem. Having preprocessed the documents of Place Reviews, we follow an automatic approach to create the UOP-dictionary.

First, we perform part-of-speech tagging in each sentence $s \in doc, \forall doc \in \mathcal{D}_R$, to classify which tag, among noun, adjective, verb, pronoun, etc., is most likely for words of s . After labeling all words, we extract a set of words considered qualifiers, i.e., marked with the adjective tag, performing a double-check using WordNet⁶, which contains a corpus of words already tagged. After that, we obtained a list with 271 commonly used words by people to qualify outdoor places. The next step is to organize these words into categories, according to their syntactic and semantic similarity, and remove the words that do not fit well in any category, being considered outliers.

To this end, we calculate a similarity score ($score_{sim}$) for each pair of words, using for that the Word2Vec model (Mikolov et al., 2013) and the sentiment polarity between words. Word2Vec is a model based on a neural network used to transform words into high-dimensional vectors that contain many linguistic regularities and patterns (Mikolov et al., 2013). It takes as input a corpus C (i.e., all sentences in documents of \mathcal{D}_R), a window size ws , a minimum count of occurrence of the word $minCount$, and a hyper-parameter m that represents the number of features. Next, the model creates a vocabulary of n unique words denoted by W from C , where each word $w \in W$ must occur at least $minCount$ times in C . The Word2Vec model computes for each pair of words $w_1, w_2 \in W : w_1 \neq w_2$, the probability to find them “nearby” into sentences of C . Two words are considered nearby if they are at most $ws - 1$ positions between them.

Using the *Places Review* dataset, where the size of *corpus* is $C = 54,612$, we trained our Word2Vec model to produce the word vectors, enabling us to identify if two different words have similar contexts, checking the similarity of

⁵ <https://www.nltk.org/howto/stem.html>. Last accessed on 12 September 2020.

⁶ <http://www.nltk.org/howto/wordnet.html>. Last accessed on 12 September 2020.

their word vectors (value denoted by $w2v_{sim}$). We empirically defined $ws = 8$, $minCount = 20$, and $m = 300$. The skip-gram model and hierarchical softmax worked better in our experiments; therefore, we employed them.

In addition to Word2Vec, we also compute the product between the sentiment polarity for each pair of words (value denoted by $sent_{sim}$), using NLTK VADER⁷ to this end. For instance, given two words, w_1 and w_2 , if both have the same sentiment polarity (negative or positive), then the $sent_{sim}$ is positive. Otherwise, the product between their sentiment polarity will be negative. Also, when one of the words has a neutral polarity, the $sent_{sim}$ becomes zero. Thus, $score_{sim}$ captures the similarity of words as follows:

$$score_{sim} = \alpha \times w2v_{sim} + (1 - \alpha) \times sent_{sim} \quad (1)$$

where $\alpha = 0.8$, since the $w2v_{sim}$ brings valuable information about the contextual relationship between words, whereas the $sent_{sim}$ works as a filter, either smoothing or increasing the $score_{sim}$ according to sentiment polarity of words.

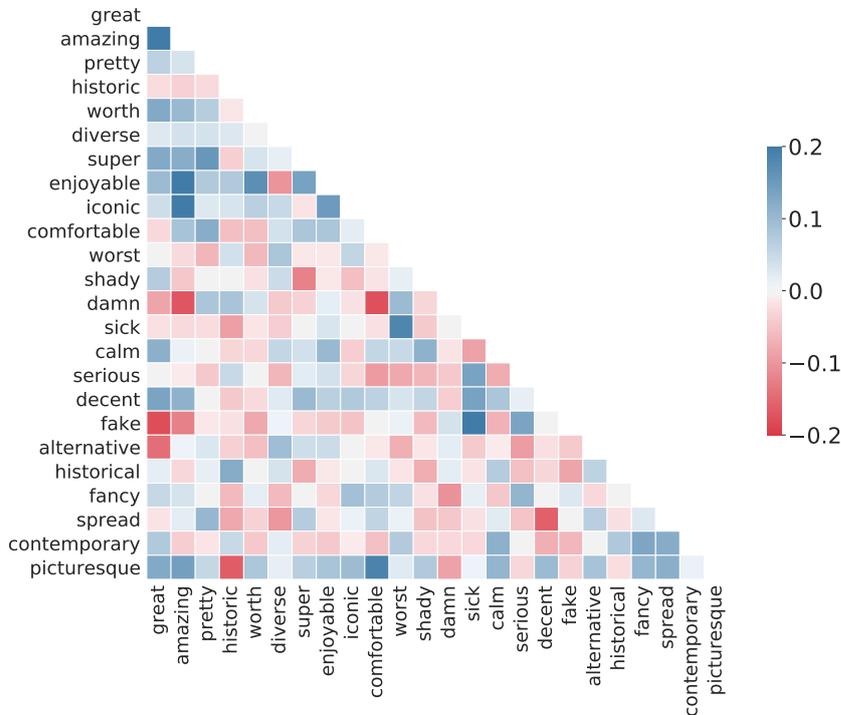


Fig. 1 Correlation matrix for some randomly selected words under consideration. [Best in color]

Figure 1 shows the $score_{sim}$ for a randomly selected sample of words from the model. As we can see, pairs of similar words, such as *amazing* and *great*,

⁷ <https://www.nltk.org/howto/sentiment.html>. Last accessed 12 September 2020.

iconic and *enjoyable*, and *damn* and *worst*, have high $score_{sim}$ value (dark blue). On the other hand, pairs of words with low $score_{sim}$ value (dark red), such as *great* and *fake*, and *damn* and *amazing*, represent dissimilar words.

In this way, we can build a graph $G = (V, E, \omega)$, where V is the set of vertices that represent the words of the model, E is the set of connections (edges) between the vertices, where each $e \in E$ has a real weight $\omega_e \in \omega$, which are defined by Eq. (1). In order to avoid meaningless connections between pairs of vertices, we calculate a threshold to eliminate “light weight” edges as follows. For each vertex $u \in V$, being E_u a set of edges incident to vertex u , i.e., $E_u = \{(u, v) : v \in V, v \neq u\}$, the threshold $thresh_u$ is defined by Eq. (2):

$$thresh_u = \overline{X_u} + \beta \times \sigma_u \quad (2)$$

where $\overline{X_u}$ denotes the weights mean of edges incident to u :

$$\overline{X_u} = \frac{\sum \omega_{u,v}}{|V| - 1} \quad (3)$$

where σ_u denotes the standard deviation:

$$\sigma_u = \sqrt{\frac{\sum (\omega_{u,v} - \overline{X_u})^2}{|V| - 2}} \quad (4)$$

and $\beta \geq 0$ defines the level of restriction. Thus, if $\beta = 0$ then $thresh_u = \overline{X_u}$, and, consequently, unmeaningful connections may remain on the graph. To avoid that, we consider a conservative value to β ($\beta = 1.13$), in order to maintain only meaningful edges. After removing from E the edges with the weight less than or equal to $thresh_u, \forall u \in V$, we obtain a graph $G' = (V', E', w')$, where $V' \subset V$ ($|V'| = 261$, after removing ten isolated vertices), $E' \subset E$ ($|E'| = 3,485$), and $w' = \{w_e : w_e \in w, \forall e \in E'\}$. With graph G' , we are able to group the vertices according to their similarity, in terms of connections.

To achieve this, we identify overlapping communities in the graph G' based on the clique percolation method (Palla et al., 2005). Thus, a k -clique community, or simply community, is the union of all cliques of size k that can be reached through adjacent (sharing $k - 1$ vertices) k -cliques (Palla et al., 2005). The clique percolation method is particularly interesting because vertices weakly connected are not placed in any community, thus removing possible noises. Also, because it allows the existence of overlapped vertices, i.e., present in multiple communities, we avoid splitting large communities into small ones with similar contexts. We empirically defined $k = 6$, which resulted in eight communities of cohesive vertices (words). These communities comprise our dictionary called UOP-dictionary.

An illustration of this dictionary is shown in Figure 2. To organize this image, we separated the communities according to their sentiment polarity. Half of the communities have words with positive polarity, and the other half has negative polarity.

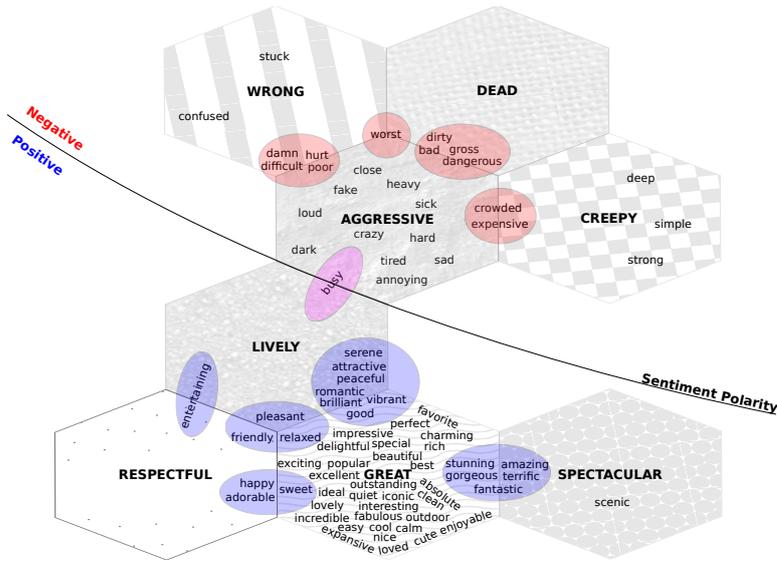


Fig. 2 Illustration of the UOP-dictionary.

Next, words of each community were placed into hexagons, where the overlapped words are highlighted with blue (positive), red (negative), or pink (transition between positive and negative) circles. To simplify the community labeling, we chose one of the available words that better represent the group’s semantics –overlapped words were not selected. This dictionary is an important step to discover the perception of outdoor places based on noisy LBSN data.

4 Extraction of Urban Perceptions

To extract the people’s perception from messages shared via LBSN about urban outdoor environments, we explored public messages (tweets) of Twitter. Twitter is an online microblogging service, where people can, among other things, share short messages of a maximum size of 280 characters. Using Twitter API, it is possible to gather tweets for areas of interest delimited by bounding boxes, where a fraction of them are geotagged (the ones we consider). Geotagged tweets have been used in many applications to predict or detect events in near real time (Kalampokis et al., 2013) and as a proxy of mobility (Lenormand et al., 2014).

To build the LBSN dataset, we collected tweets, from January to August 2018, for Chicago, NYC, and London. We also represent the LBSN dataset as a document collection defined as:

Definition 2 Collection \mathcal{D}_L , where each $doc \in \mathcal{D}_L$ is a document determined by a tuple $doc = (id, s, \tau, g)$, where id is a unique identifier, s is a list of sen-

tences that comprise all content written by the user, $\tau \in \mathbb{R}$ is a timestamp, and $g \in \mathbb{R} \times \mathbb{R}$ is the geographic coordinates, expressed by latitude and longitude, respectively.

When applying the UOP-dictionary in \mathcal{D}_L , several documents not related to individual perception about urban outdoor areas tend to be retrieved. The reason for that is because the qualifiers that compose the dictionary are not restricted to describe places. Some of them can also be used to describe other entities, such as people and things. To address this problem, we group documents that have both spatial-temporal and semantic similarities, and, thus, disregard documents non-related to urban areas, or individual perceptions unrelated to urban outdoor areas. To this end, we propose an unsupervised clustering algorithm to group documents with *spatial*, *temporal*, and *semantic* similarities present in the data.

First, we remove any spatial noise from the dataset containing LBSN data, called \mathcal{D}_L , i.e., all documents whose geolocation is coincident to others according to a threshold, $thresh_{spatial}$, are filtered out from \mathcal{D}_L . The $thresh_{spatial}$ indicates the acceptable maximum number of documents with a coinciding GPS location. The probability of this situation to happen in practice is low for real data, being more common in automatic processes, such as those used by web robots, as discussed in (Tasse et al., 2017). Thus, this step is essential to avoid considering invalid data. Based on the documents \mathcal{D}_L , we have determined a threshold, $thresh_{spatial} = 10$, to perform this filtering since most geolocations (about 90%) have less than 10 documents associated with them. This process produces the collection \mathcal{D}'_L .

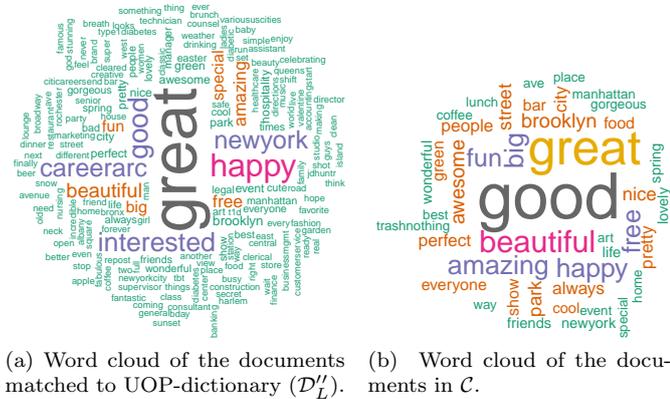
Next, we use the UOP-dictionary to label the appropriate category of perception of documents in \mathcal{D}'_L .

In this way, a given document with at least one word corresponding to the UOP-dictionary can be labeled with one of the eight labels of the dictionary. If two or more words are present in the dictionary, multiple labels can be assigned to the same document. This helps to reflect the heterogeneity of perceptions, where the same area may have distinct perceptions according to the people’s opinion. Alternatively, the document may not have words matching with the dictionary, which implies that it may not be related to an outdoor area qualification and, therefore, is disregarded. As a result, we have a new collection of documents \mathcal{D}_L'' , where the documents are properly labeled with at least one community label of the UOP-dictionary.

After that, we perform a semantic similarity task. For each document $doc \in \mathcal{D}_L''$, we explore the Word2Vec model to calculate the likelihood of a doc be a member of a specific community of the UOP-dictionary (Taddy, 2015). This procedure results in a score, ranging from 0, very unrelated, to 100, very similar, and we explore it to decide if the doc has enough semantic similarity with the model. Based on the dataset \mathcal{D}_L'' , we have determined a threshold, $thresh_{semantic} = 18$, to perform this classification, where most documents (about 75%) have a score higher than 18. To define this value, we studied the distribution of scores. This enabled us to observe a clear change on the curve

Table 2 Number of documents after main steps of our approach.

	Chicago	London	NYC
Total of documents (\mathcal{D}_L)	803,689	1,364,922	2,690,327
# documents after applying the spatial filter (\mathcal{D}'_L)	181,836	345,100	620,350
# documents match with the UOP-dictionary (\mathcal{D}''_L)	24,951	58,482	73,062
# documents after applying the semantic filter (\mathcal{D}'''_L)	18,293	39,760	54,539

**Fig. 3** Word clouds before and after applying our approach.

when the score is 18, indicating that after that point, documents tend to be more related to the context of interest. In fact, when evaluating documents with score values less than 18, we start to find documents unrelated to urban outdoor areas. On the other hand, for scores slightly above 18, we practically find only documents related to the urban outdoor context.

Table 2 summarizes the number of documents remaining after each step of our approach. The last line highlights the total of documents in \mathcal{D}'''_L , representing semantically similar documents, which are clustered according to their spatial-temporal similarity. To this end, we split the documents monthly, according to the time they were shared. Also, we used the Hierarchical Density-Based Spatial Clustering of Application with Noise (HDBSCAN) (Campello et al., 2013), for clustering the documents based on their spatial similarity. HDBSCAN extends the DBSCAN (Ester et al., 1996), using a technique to find a clustering that gives the best stability over an ϵ value (Campello et al., 2013). To explore HDBSCAN for this task is interesting since it is based on a neighborhood density to identify spatial clusters. Also, it can create clusters with different formats and sizes, and, typically, achieves satisfactory results even in the presence of noises. This clustering process results in a set of clusters \mathcal{C} .

To illustrate the potential of our approach to extract urban perceptions, Figure 3 shows two word clouds, where more centralized words and with larger font sizes are the most frequent ones, considering data collected from NYC. In Figure 3(a), before applying our approach, we can observe that several frequent words are not related to urban outdoor areas, e.g., “careerarc” and

“interested”. On the other hand, the remaining documents, i.e., the documents in \mathcal{C} , are strongly related to urban outdoor perceptions as displayed in Figure 3(b), being more suitable to be used on perception mapping.

5 Experiments and Evaluation

This section discusses the results and evaluation of our proposed approach, considering Chicago, New York City (NYC) and Greater London as our scenarios (see Appendix A). For each studied city, we selected three neighborhoods with the objective of evaluating diverse areas in each city. To this end, we selected representatives of rich/central areas, problematic areas (e.g., poor and/or violent), and residential and non-problematic areas, favoring neighborhoods with more data.

5.1 General Map at City Scale

With the help of interactive maps built for this study, urban perceptions are mapped into the cities according to the number of observed collective perceptions. This tool eases the process of studying the perceptions found. For instance, Figure 4 shows the heat maps of clusters for each perception category separately, considering all data on the evaluated areas from January to August 2018 (see Appendix B for other cities). As we can see, there is an interactive selector to filter the desired perceptions, which help to analyze each perception individually. Besides, our application visually combines close clusters into one single representation (circle) and simplifies the display of clusters on the map. The number on a circle indicates how many points in all grouped clusters it contains – colors from green (low) to red (high) indicate the number of points. As the zoom increases, the clusters are split, and individual markers will, eventually, appear. Such maps are useful to show the perception concentration at the city scale.

5.2 Spatial Analysis at Neighborhood Scale

In this section, we present the main characteristics of urban areas evaluated and map the spatial distribution of perceptions using markers, with “Plus/Minus” icons to represent the positive/negative perceptions. We also assigned colors to distinguish them: dark blue (GREAT); gray (LIVELY); blue (RESPECTFUL); light blue (SPECTACULAR); dark red (AGGRESSIVE); red (WRONG); light red (DEAD); and beige (CREEPY).

The areas studied in Chicago are: Loop, Norwood Park and Wicker Park. Loop is the downtown area of Chicago, known as an important commercial and financial center of the city, attracting crowds of visitors and residents with different profiles to their several types of venues.

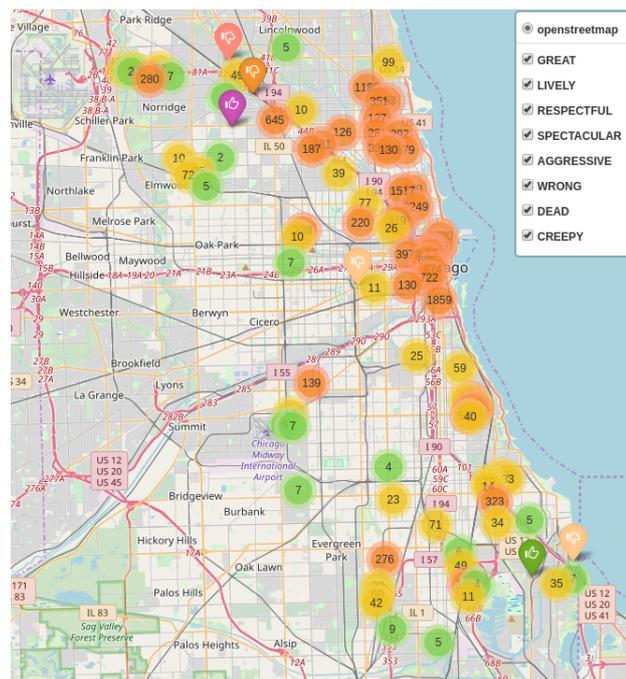


Fig. 4 Screenshot of an interactive urban perception map in Chicago. This tool was built for this study to help to evaluate the perceptions found.

As we can see in Figure 5, Chicago's Downtown area contains a considerable concentration of distinct perceptions, suggesting that most of them coexist in the neighborhood. Such finding is not surprising because Downtown areas tend to present a wide variety of sounds, visual elements, odors, among other characteristics, which can potentially cause distinct perceptions to people. For this reason, different urban perceptions may be favored to occur.

Norwood Park is mostly residential, with some leisure area options, such as parks with sports fields and paths, and a hospital near to the central area of the neighborhood. As shown in Figure 6, there are *AGGRESSIVE* points spread in different parts of the neighborhood, which indicate that some people had bad experiences in this region. Beyond this perception, only *GREAT*, *LIVELY*, and *RESPECTFUL* perceptions have occurred in the neighborhood during the considered period. Note also that we have comparatively fewer data concerning the other areas. Nevertheless, the proposed approach is still able to gather overall perceptions about this area.

Considering the Wicker Park region, known as a hub for shopping, eating, and cultural activities in the city, we can see in Figure 7 several overlaps among the perceptions in the intersection between the "N Milwaukee" and "W North" avenues, where restaurants and bars are concentrated. This indicates that, for example, shopping or eating in these areas is an everyday activity to be performed, attracting a large number of citizens and visitors, inducing

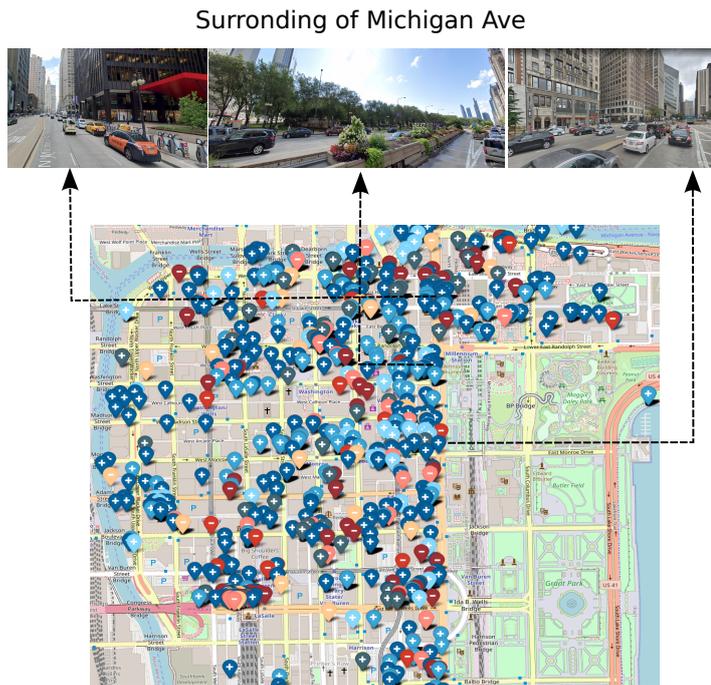


Fig. 5 Perceptions map in Loop (Downtown). [Best in color]

a mix of perceptions. We can also see that another interesting finding in this neighborhood is located close to Dean Playground Park. In this spot, there are cultural venues, favoring positive perceptions.



Fig. 6 Perceptions map in Norwood Park. [Best in color]



Fig. 7 Perceptions map in Wicker Park. [Best in color]

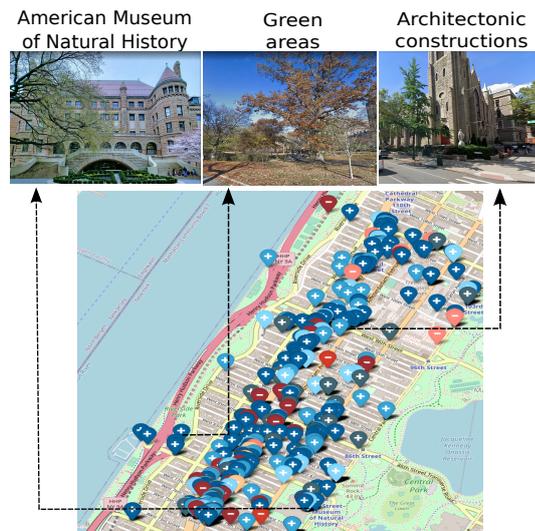


Fig. 8 Perceptions map in Upper West Side. [Best in color]

Studying now NYC, the three regions in our evaluation are: Upper West Side, Park Slope and Gowanus, and Jamaica. Upper West Side is famous for the architectonic constructions. Besides, it offers several cultural venues, such as the American Museum of Natural History. It also concentrates an academic community due to its proximity to Columbia University, and has several options for gastronomy and nightlife. The perceptions in the Upper West Side, Figure 8, are concentrated around the Broadway and Amsterdam Avenues, and also spread by the neighborhood close to the shopping and green areas. The mainly perceptions are *GREAT*, *SPECTACULAR* and *RESPECTFUL*.

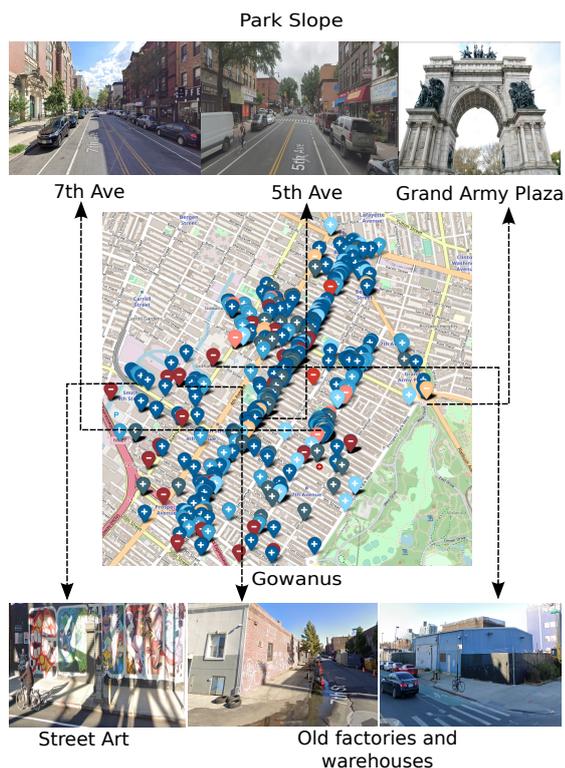


Fig. 9 Perceptions map in Park Slope and Gowanus. [Best in color]

The Park Slope and Gowanus are located in Brooklyn and, despite the close geographic proximity, they have very distinct characteristics, as shown in Figure 9. Park Slope is famous for the charming buildings, great food establishments and lovely parks, such as Prospect Park and Grand Arm Plaza, which favor to emerge several positive perceptions, mainly the surrounding of 5th Avenue, Union Street, Barclays Center and along 7th Avenue. On the other hand, Gowanus is an industrial business zone with a few residential areas mixed among old factories and warehouses, where there are several points of perception *AGGRESSIVE*. Recently, Gowanus has been attracting artists to

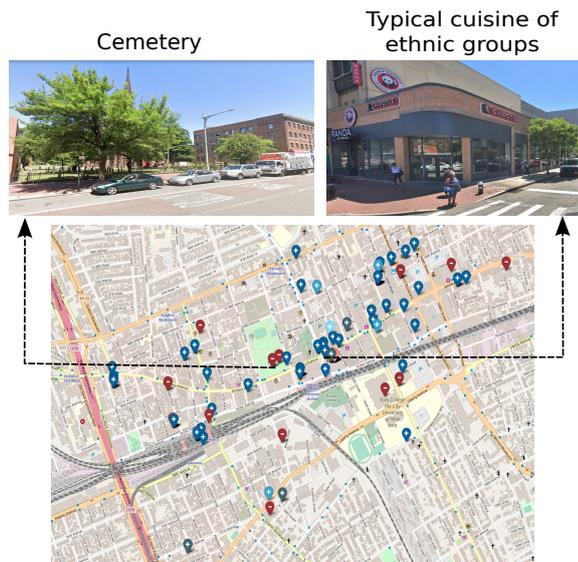


Fig. 10 Perceptions map in Jamaica. [Best in color]

the neighborhood, who have made a significant change to the area, such as street art murals and vibrant bars.

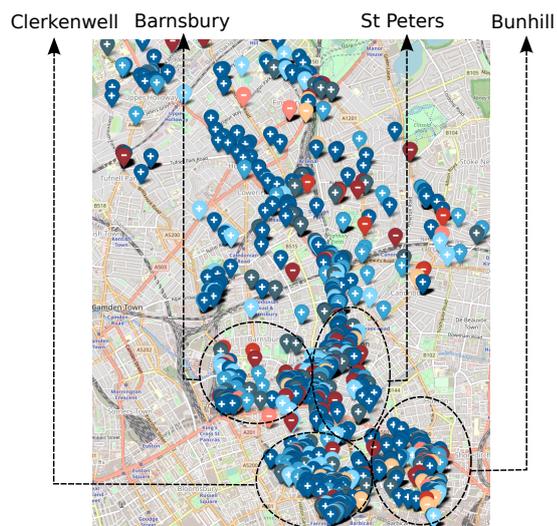


Fig. 11 Perceptions map in Islington. [Best in color]

Situated in the Queens borough, Jamaica is a residential neighborhood that is mainly populated with Latinos (Hinterland et al., 2018). Other minority populations, such as Asians and African Americans, also live in large

numbers in this area (Hinterland et al., 2018). Due to the high diversity of the community living in the neighborhood, there are many restaurants and grocery stores with typical cuisine of ethnic groups, arising positive perceptions *GREAT*, *LIVELY* and *RESPECTFUL* close to those places, as shown in Figure 10. By contrast, Jamaica suffers from a high rate of violent crimes (Hinterland et al., 2018), being a negative aspect of the neighborhood, where several points of *AGGRESSIVE* are spread in the area.

Finally, we evaluated London Boroughs rather than neighborhoods, since most neighborhoods have small areas, where most of them have insufficient data to perform the analysis. The selected regions for evaluating are Islington, Hammersmith and Fulham, and Lambeth. Islington concentrates diverse perception across its area, as shown in Figure 11. In that figure, we can also see that the positive ones stand out. Among Islington districts, we can highlight Bunhill, Clerkenwell, Barnsbury, and Saint Peters, as the districts that concentrate the highest number of perceptions. This fact might be motivated by the proximity of these districts with London’s downtown, i.e., which implies a higher diversity and a higher number of people.

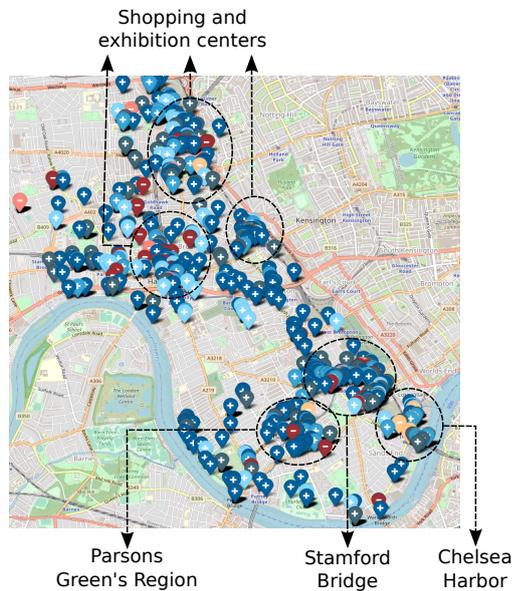


Fig. 12 Perceptions map in Hammersmith and Fulham. [Best in color]

With one of the most expensive places to live in London, Hammersmith and Fulham have been attracting affluent families and prosperous young people to live there. Moreover, Hammersmith and Fulham house old-fashion pubs, soccer stadiums, cultural venues, shops, cafes, and malls, being a popular place to visit in the city. In Hammersmith and Fulham, we can see in Figure 12, close to Stamford Bridge, several points of positive perceptions, mainly *GREAT*,

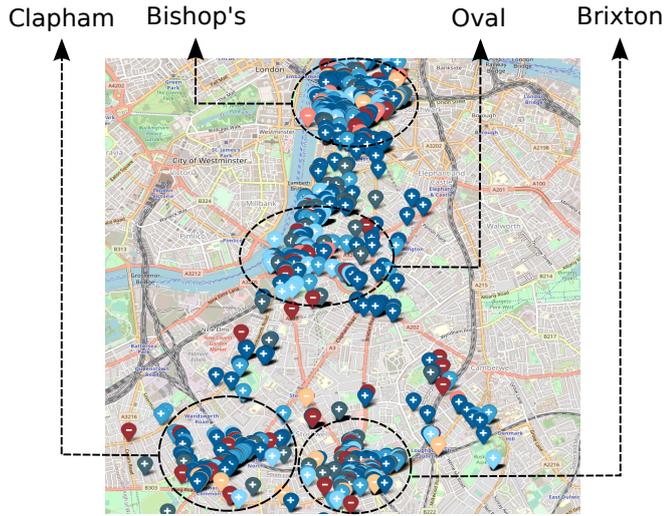


Fig. 13 Perceptions map in Lambeth. [Best in color]

LIVELY, and *RESPECTFUL* categories, mixed with points of the perception *AGGRESSIVE*, the negative one. Moreover, there are dense groups of positive category points in various parts of the borough. For instance, in Parsons Green and Chelsea Harbour in Fulham, and Shopping centers and exhibition centers in Hammersmith, indicating that these areas are desirable to visit. We can also see some points of categories *AGGRESSIVE* and *CREEPY* close by and, in some cases, overlapping those groups, suggesting, probably, problems inherent to crowded places.

Turning our attention to Lambeth, as we can see in Figure 13, there is a concentration of perceptions in the following districts: Bishop's, Oval, Brixton, and Clapham. Besides, there are several groups of mixed perceptions spread in the remaining areas of Lambeth. Despite common unbalance in the number of perceptions in other scenarios, mainly with *GREAT* perception, almost all perceptions appear in a significant quantity in this region.

5.3 Temporal Analysis at Neighborhood Scale

To evaluate the urban perceptions extracted from the studied areas, we conducted a temporal analysis to measure the persistence of perceptions identified over time. Such analysis is useful for ranking the perceptions in the neighborhoods, as well as to identify if perceptions keep unchanged independent of the period. To this end, we calculate the z-scores similar to some authors, such as (Aiello et al., 2016; Quercia et al., 2016; Redi et al., 2018), defined as follow:

$$\text{z-score}_i^j(n) = \frac{X_i^j(n) - \mu(X_i^j)}{\sigma(X_i^j)}, \quad (5)$$

where i denotes one specific perception category among all we consider, j denotes the period (here we consider months), and n denotes a specific city neighborhood among all neighborhoods. In addition, $X_i^j(n)$ indicates the number of points of perception i in the neighborhood n during the month j , whereas $\mu(X_i^j)$ and $\sigma(X_i^j)$ are the mean number and the standard deviation of points of perception i during the month j in every neighborhood of the city, respectively. Using z-score, we can capture, for every neighborhood, the perceptions that stand out, even when the number of samples of perception categories is very unbalanced, which is our case. For instance, the category *GREAT* is by far the most numerous, and the reason for that could be diverse, e.g., people tend to share more the “good” moments in social media.

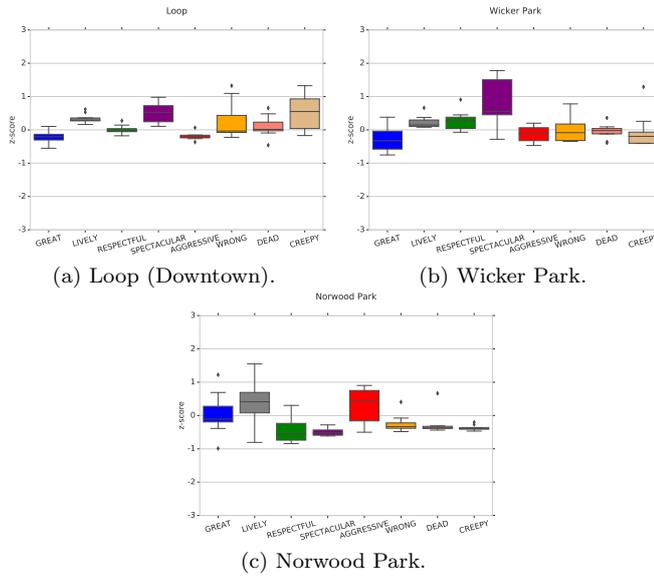


Fig. 14 Boxplots containing the perception strength of Chicago’s neighborhoods.

Figure 14 shows the results (exploring boxplots) for Chicago. As we can see in Figure 14(a), several perceptions occurred in the Downtown area with low variation along the time (less than one standard deviation in most cases) and similar intensity (median z-score near to 0). Among them, *LIVELY*, *SPECTACULAR*, and *CREEPLY* are the most significant ones and tend to occur with high strength in this area. *WRONG* can appear with high intensity on few months ($z\text{-score} \geq 1$); however, note that the median is close to 0. This is not the case for *CREEPLY* perception, indicating that visitors recurrently had some issues during their personal experience when visiting the neighborhood, generating this class of perception.

Wicker Park is predominantly *SPECTACULAR* with the higher median perception strength, as shown in Figure 14(b). Moreover, *LIVELY* and *RE-*

SPECTFUL also have a significant occurrence in the neighborhood. At the same time, the remaining perceptions keep the median perception strength below the expected to a neighborhood (z -score < 0). The perception about the region of Norwood Park neighborhood is predominantly *LIVELY*, followed by *AGGRESSIVE*. Note that for other perceptions, the median strength is below 0, indicating they are hardly related to this area.

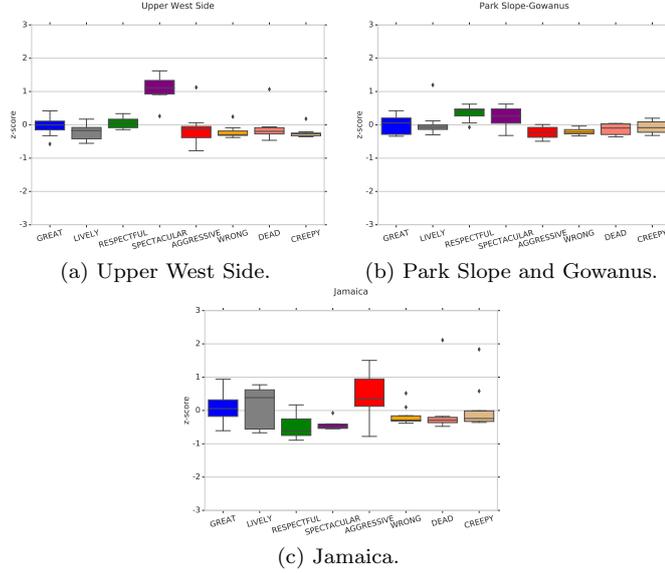


Fig. 15 Boxplots containing the perception strength of NYC's neighborhoods.

Turning our attention to NYC, Figure 15 shows the results for its studied neighborhoods. As shown in Figure 15(a), the Upper West Side region is mainly *SPECTACULAR* (median z -score ≥ 1), but also is *GREAT* and *RESPECTFUL*. Negative perceptions have low strength overall. Studying Park Slope and Gowanus neighborhoods, Figure 15(b), most perceptions kept stable throughout the months. *GREAT*, *RESPECTFUL* and *SPECTACULAR* are the most intense ones in the area. For this area, negative perceptions are also not significant (median z -score < 0). As we can see in Figure 15(c), the median perception strength of *LIVELY* and *AGGRESSIVE* is the most predominant in Jamaica neighborhood. As we can see, *GREAT* tends to occur moderately in this region.

Finally, Figure 16 shows the urban perceptions for London. It is possible to see that Hammersmith and Fulham (Figure 16(a)) has *LIVELY* and *SPECTACULAR* as predominant perceptions. Note that the perception *RESPECTFUL* could be considered a secondary perception to this region, with the z -score value close to zero and low variation. Other perceptions have low z -score values, mainly the negative ones, being perceptions less relevant. For Is-

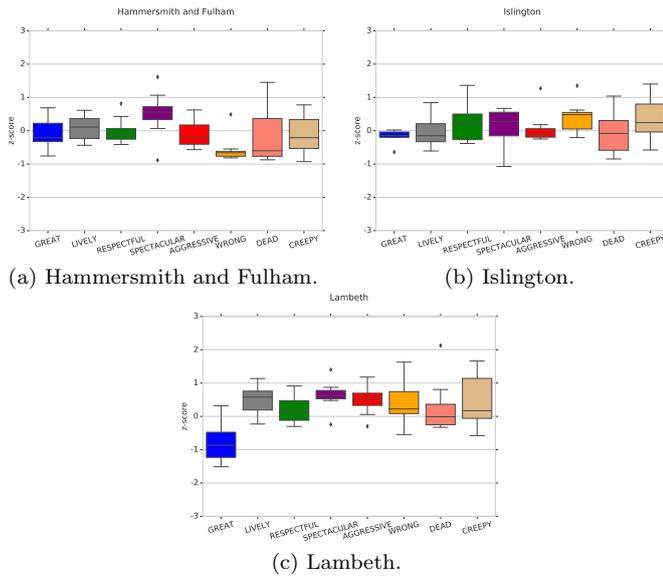


Fig. 16 Boxplots containing the perception strength of London’s neighborhoods.

lington, Figure 16(b), we observe an equilibrium among positive and negative perceptions. *RESPECTFUL* and *SPECTACULAR* are the most significant positive perceptions, and *WRONG* and *CREEPY* are the negative perceptions more notable. Lambeth borough shows the most heterogeneous behavior among the studied areas (Figure 16(c)), where six of the eight perceptions captured in this area have high z -score values. We can see that *LIVELY* and *SPECTACULAR*, followed by *AGGRESSIVE* and *CREEPY*, are the most relevant in this area.

5.4 Comparison with Place Pulse 2.0

To verify if the perception identified by our method reflects the overall people’s opinion, we have conducted a comparative analysis with results obtained by the research project Place Pulse 2.0 (Dubey et al., 2016). This project collected volunteers’ perception based on features present in urban outdoor images, containing more than 1,2M comparisons between pairs of images. According to volunteers’ opinions, the comparison captures which image better matches with a particular perception regarding safety, wealthy, beautiful, lively, boring, and depressing. More information about the process can be obtained in (Dubey et al., 2016). The researchers collected opinions based on images of 56 cities, among them Chicago, NYC, and London (responses used in this study).

Since there is no direct mapping between our perception classes and the ones considered by Place Pulse, we aggregated them according to their sentiment polarity. This was done based on the description given by the authors

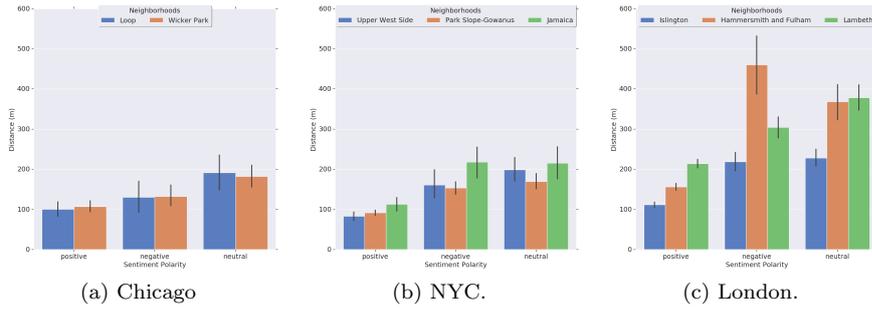


Fig. 17 The average distance between the nearest points from Place Pulse 2.0 data (Dubey et al., 2016) and our approach, according to their polarity in each city and neighborhood. [Best in color]

of the project (Dubey et al., 2016). Thus, we considered the Place Pulse categories wealthy, beautiful, and safety (because they represent places that look safer) as positive. Boring and depressing as negative. Moreover, the “lively” category as neutral, since it appears between positive and negative categories in our approach (see Figure 2). In this way, we performed a comparative analysis of perceptions extracted via our approach and Place Pulse, despite the different labels used in these approaches.

First, we analyzed the spatial similarity between perception points from both approaches to understand how near they are. We do that as follows. Given a point p with a particular polarity according to Place Pulse, we find among points with the same polarity extracted via our approach the one whose distance to p is minimum. We do this process for all points from Place Pulse. According to this process, Figure 17 shows the average distance (with 95% confidence interval) from all points, separated according to sentiment class and neighborhood for each city. As there are no points with perceptions in the Norwood Park neighborhood on Place Pulse data, we ignored this region in forthcoming results.

As we can see, the average distance follows a similar behavior in most cases. Positive images tend to have the shortest distances, followed by negative and neutral ones. One of the most prominent exceptions is Hammersmith and Fulham boroughs, where points with negative polarity have the most significant average distance compared to others. The average distance in all cases is around 100 to 200 meters for American cities. For London, the average distance range is higher, about 100 to 400 meters. This is comprehensible since the studied areas in London are considerably larger than the others. By this analysis, we observed that, in general, there are few spatial divergences between our approach and Place Pulse. This suggests that both approaches tend to agree reasonably well on the sentiment reflected by urban areas.

After that, we evaluated the perception strength using Place Pulse data based on z -score, as defined in Eq. 5. We aggregated the perceptions according to their sentiment polarity to facilitate the comparison. Figure 18 shows the results for the evaluated areas of Chicago. According to Place Pulse data,

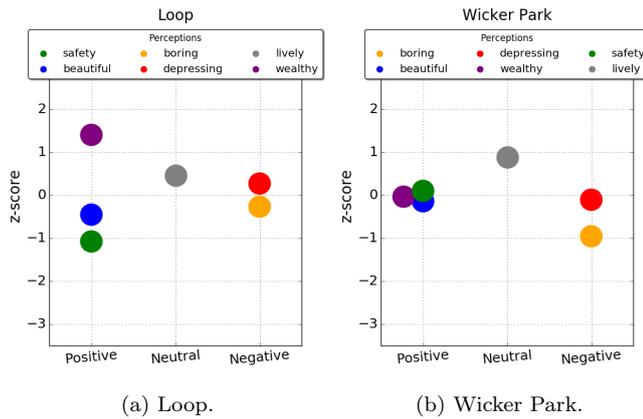


Fig. 18 The perception strength of Chicago's neighborhoods according to Place Pulse 2.0 data (Dubey et al., 2016). [Best in color]

Loop is mainly *wealthy*, but is also *lively* and *depressing*, with less intensity as shown in Figure 18(a). The similarity between the perception identified by our approach with this result is striking, where the most intense perception is also positive (*SPECTACULAR*) followed by *LIVELY* and negative perceptions. The perception about Wicker Park, Figure 18(b), is predominantly *lively*. Note that the positive perceptions are concentrated close to zero, being the safety perception the most significant among them. The negative perceptions also have less relevance in this region. According to our results, the *LIVELY* category also stands out (third most prominent). At the same time, the *SPECTACULAR* and *RESPECTFUL* are the primary and secondary categories, respectively, with the negative perceptions keeping the average z -score below zero.

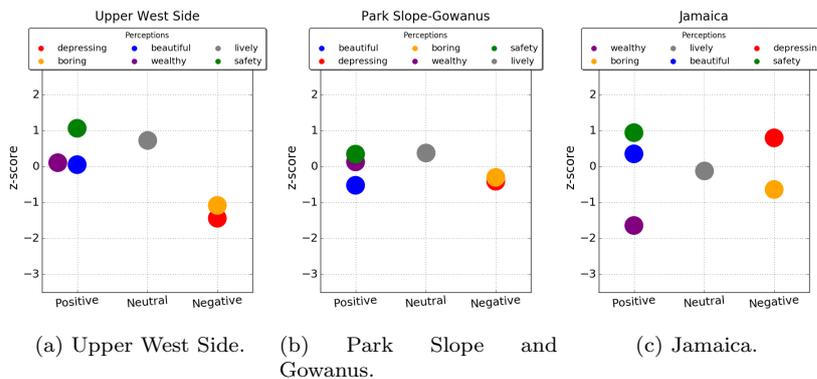


Fig. 19 The perception strength of NYC's neighborhoods according to Place Pulse 2.0 data (Dubey et al., 2016). [Best in color]

Similarly, Figure 19 shows the results for evaluated areas of NYC. As shown in Figure 19(a), Upper West Side according to the Place Pulse data is mostly *safety* and *lively*, followed by *wealthy* and *beautiful*. Besides, the negative perceptions have low strength in this region (z -score < -1). According to our results, positive perceptions also stand with strength close to zero, except category *SPECTACULAR* that exceeds the others. Our results also identified low strength for negative perceptions in this area. Only the neutral category shows a slight difference, where *LIVELY* is the fourth most evident perception for such region. Nevertheless, the results show a strong correlation between the approaches. For Park Slope and Gowanus, Figure 19(b), all perceptions keep the z -score close to zero. Thus, it is an area where no perception stands out. Note that our results have a significant match for the same area, following a similar behavior. The results for the Jamaica neighborhood, Figure 19(c), show a high intensity of *depressing* perception, being the second most prominent perception in this area according to Place Pulse. Besides, *safety* and *beautiful* perceptions are the most relevant positive categories, while the *lively* perception has z -score slight below zero. According to our results, a negative perception, *AGGRESSIVE*, also has high strength in this area, followed by *LIVELY* and *GREAT*. Despite the divergence on the rank order of perceptions, the overall picture is still well captured.

Finally, Figure 20 shows the results for London. According to Place Pulse, the perception about Hammersmith and Fulham is predominantly *wealthy* and *safety*, followed by *depressing* and *lively*, as shown in Figure 20(a). Our approach also identified a very similar perception for this area, where the *SPECTACULAR* category corresponds to a primary perception of the neighborhood. At the same time, *LIVELY* is a secondary category. Furthermore, our results show that the *AGGRESSIVE* category has a similar strength to other positive perceptions. This indicates that it might coexist, as the case of the *depressing* perception, according to Place Pulse in this area. For Islington borough, Figure 20(b), the positive and neutral perceptions have a good match between our results and Place Pulse. However, according to our results, this area also suffers the influence of negative perceptions, mainly the *WRONG* and *CREEPY* categories. In contrast, Place Pulse results show low relevance of negative perceptions in the same region. Similarly, our results show a high intensity of several positive, neutral and negative perceptions at Lambeth.

In general, our approach, compared with the Pulse Place data, achieved a satisfactory similarity between extracted perception by both strategies in most cases.

6 Discussions and Limitations

Our work has the potential to help people to extract knowledge from the city, and, thus, to improve its understanding. This is useful for many tasks, for instance, to assist in the development of intelligent services, such as personalized route recommendations, which could be offered by systems like Waze

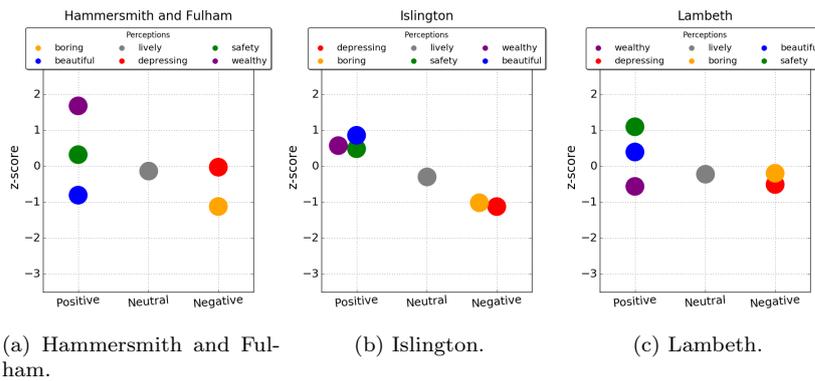


Fig. 20 The perception strength of London’s neighborhoods according to Place Pulse 2.0 data (Dubey et al., 2016). [Best in color]

or Google Maps. In this context, tourists can explicitly request routes to walk through areas with selected perceptions.

It is also important to note that our approach could be applied, with the proper adaptations, to other domains as well. In this study, we focused on urban outdoor areas, but our approach could be explored, for example, for indoor spaces. By employing our methodology in other review datasets, perhaps larger and more updated than those explored, we could potentially enrich our dictionary with new significant words. This is important because the language is in constant evolution and can change in unique ways in different locations.

Place Pulse 2.0 considers visual patches of images to identify the place’s urban perception instead of free-texts shared by individuals while visiting the area, which can potentially differ the people’s perception since that non-visual features can influence in their perception. However, it is one of the few studies that mapped the urban perception in a large scale, considering a significant number of categories, and provided their data, giving us a piece of general knowledge about urban perception of outdoor areas.

We are also aware of some limitations of our proposal. Upon analyzing samples of messages shared by individuals, we found certain phrases containing conflicting perceptions. For example, “... *What an energetic, **amazing**, fun, **loud**, show*” and “... *For anyone having a **bad** week or needs a **happy** break, here’s a polar bear ...*”. In those examples, positive and negative perceptions would be identified. Our approach neutralized this problem with the proposed clustering step; however, if this type of situation happens many times in the same area, this could favor identifying wrongly two types of perceptions, depending on the values for the approach’s parameters. We did not find any relevant problem regarding these cases in our results. However, the application of our approach in other datasets might face this possible limitation.

Besides, our approach takes into account the opinions of Twitter users who shared their perceptions when visiting the evaluated areas. Despite the clear

advantages of using social media data to extract the overall picture of urban areas' perceptions, such findings might not necessarily correspond to accurate truth. This is because just a few population groups are more likely to use social media, especially adults (18+), urban dwellers, and high-income people (Perrin and Anderson, 2019). Moreover, the human perception spectrum arising from the data sources considered in our study can be less expressive than those that could have been obtained by traditional surveys. However, note that other data sources, perhaps more expressive towards specific perceptions, could also be explored by our methodology, and thus, being complementary instead of replacing them.

7 Conclusion

In this study, we presented an automatic approach to support the learning and mapping of perceptions of urban outdoor areas from an extensive collection of noisy data expressing individuals' opinions in LBSNs. Due to several advantages of our approach, e.g., easier scalability, it has the potential to complement traditional methods, such as surveys. In this sense, scenarios where traditional surveys have been conducted, our approach can take advantage of their data (assuming availability) to learn about citizens' perceptions and proceed with our methodology's steps. Also, in scenarios where there is no information about urban perceptions, our approach can support bring such information, helping people better understand the semantics existing in different city regions. As future work, we intend to evolve this approach to incorporate other data sources, such as Instagram and Facebook. Besides, we plan to apply and evaluate our strategy for content in different languages and analyze more details about the implications of other variables, such as weather, on the results.

Acknowledgements This work was partially supported by Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001 (process 88881.337131/2019-01), project URBCOMP (Grant #403260/2016-7 from National Council for Scientific and Technological Development agency - CNPq), and GoodWeb, BigCloud and Mobilis (Grants #2018/23011-1, #2015/24494-8 & #2018/23064-8 from São Paulo Research Foundation - FAPESP).

References

- Aiello LM, Schifanella R, Quercia D, Aletta F (2016) Chatty maps: constructing sound maps of urban areas from social media data. *Open Science* 3(3):150690
- Campello RJ, Moulavi D, Sander J (2013) Density-based clustering based on hierarchical density estimates. In: *Proc. of PAKDD*, Springer, Gold Coast, Australia, pp 160–172

- De Nadai M, Vieriu RL, Zen G, Dragicevic S, Naik N, Caraviello M, Hidalgo CA, Sebe N, Lepri B (2016) Are safer looking neighborhoods more lively? a multimodal investigation into urban life. In: Proceedings of the 24th ACM international conference on Multimedia, pp 1127–1135
- Dubey A, Naik N, Parikh D, Raskar R, Hidalgo CA (2016) Deep learning the city: Quantifying urban perception at a global scale. In: Proc. of ECCV, Springer, pp 196–212
- Ester M, Kriegel HP, Sander J, Xu X, et al. (1996) A density-based algorithm for discovering clusters in large spatial databases with noise. In: Proc. of KDD, AAAI Press, Portland, United States, pp 226–231
- Henshaw V (2013) Urban smellscapes: Understanding and designing city smell environments. Routledge
- Hinterland K, Naidoo M, King L, Lewin V, Myerson G, Noubissi B, Woodward M, Gould LH, Gwynn RC, Barbot O, Bassett MT (2018) Community Health Profiles 2018, Queens Community District 12: Jamaica and Hollis. NYC.gov, URL <https://www1.nyc.gov/assets/doh/downloads/pdf/data/2018chp-qn12.pdf>
- Hsu YC, Cross J, Dille P, Tasota M, Dias B, Sargent R, Huang THK, Nourbakhsh I (2019) Smell pittsburgh: community-empowered mobile smell reporting system. In: Proc. of IUI, ACM, Los Angeles, USA, pp 65–79
- Hu Y, Deng C, Zhou Z (2019) A semantic and sentiment analysis on online neighborhood reviews for understanding the perceptions of people toward their living environments. *Annals of the American Association of Geographers* 109(4):1052–1073
- Jenkins A, Croitoru A, Crooks AT, Stefanidis A (2016) Crowdsourcing a collective sense of place. *PloS one* 11(4):e0152932
- Jiang S, Qian X, Mei T, Fu Y (2016) Personalized travel sequence recommendation on multi-source big social media. *IEEE Transactions on Big Data* 2(1):43–56
- Kalampokis E, Tambouris E, Tarabanis K (2013) Understanding the predictive power of social media. *Internet Research* 23(5):544–559
- Keizer K, Lindenberg S, Steg L (2008) The spreading of disorder. *Science* 322(5908):1681–1685
- Lenormand M, Picornell M, Cantú-Ros OG, Tugores A, Louail T, Herranz R, Barthelemy M, Frias-Martinez E, Ramasco JJ (2014) Cross-checking different sources of mobility information. *PloS one* 9(8):e105184
- Mikolov T, Sutskever I, Chen K, Corrado GS, Dean J (2013) Distributed representations of words and phrases and their compositionality. *Adv in Neural Inf Proc Sys* pp 3111–3119
- Naik N, Philipoom J, Raskar R, Hidalgo C (2014) Streetscore-predicting the perceived safety of one million streetscapes. In: Proc. of CVPR Workshops, Columbus, USA, pp 779–785
- Nasar JL (1990) The evaluative image of the city. *Journal of the American Planning Association* 56(1):41–53
- Palla G, Derényi I, Farkas I, Vicsek T (2005) Uncovering the overlapping community structure of complex networks in nature and society. *nature*

- 435(7043):814
- Pareek U (1988) *Organizational Behaviour Processes*. Rawat Publications, Jaipur, URL <https://books.google.com.br/books?id=hZdHAgAACAAJ>
- Perrin A, Anderson M (2019) Share of U.S. adults using social media, including Facebook, is mostly unchanged since 2018. Pew Research Center, URL <http://tiny.cc/9nlxsz>
- Porzi L, Rota Bulò S, Lepri B, Ricci E (2015) Predicting and understanding urban perception with convolutional neural networks. In: Proc. of MM, ACM, Brisbane, Australia, pp 139–148
- Quercia D, O’Hare NK, Cramer H (2014) Aesthetic capital: what makes london look beautiful, quiet, and happy? In: Proc. of CSCW, ACM, Vancouver, Canada, pp 945–955
- Quercia D, Schifanella R, Aiello LM, Kate M, et al. (2015) Smelly maps: The digital life of urban smellscapes. In: Proc. of ICWSM, AAAI Press, Oxford, UK, pp 327–336
- Quercia D, Aiello LM, Schifanella R (2016) The emotional and chromatic layers of urban smells. In: Proc. of ICWSM, AAAI Press, Cologne, Germany
- Redi M, Aiello LM, Schifanella R, Quercia D (2018) The spirit of the city: Using social media to capture neighborhood ambiance. Proc of CHI 2:144
- Robbins S (2003) *Organizational Behavior*. Prentice-Hall International editions, Prentice Hall
- Ross CE, Mirowsky J (2001) Neighborhood disadvantage, disorder, and health. *Jour of health and social beh* pp 258–276
- Santani D, Ruiz-Correa S, Gatica-Perez D (2018) Looking south: Learning urban perception in developing cities. *ACM Transactions on Social Computing* 1(3):13
- Shaban H (2019) Twitter reveals its daily active user numbers for the first time. *Washington Post*, URL <https://www.washingtonpost.com/technology/2019/02/07/twitter-reveals-its-daily-active-user-numbers-first-time/>
- Smith A, Anderson M (2018) Social media use in 2018. URL <https://www.pewinternet.org/2018/03/01/social-media-use-in-2018>
- Steiger E, Resch B, Zipf A (2016) Exploration of spatiotemporal and semantic clusters of twitter data using unsupervised neural networks. *Jour of Geo Inf Sci* 30(9):1694–1716
- Taddy M (2015) Document classification by inversion of distributed language representations. In: Proc. of ACL, Beijing, China, vol 2, pp 45–49
- Tasse D, Liu Z, Sciuto A, Hong JI (2017) State of the geotags: Motivations and recent changes. In: Proc. of ICWSM, Montreal, Canada, pp 250–259
- Yao Y, Li X, Liu X, Liu P, Liang Z, Zhang J, Mai K (2017) Sensing spatial distribution of urban land use by integrating points-of-interest and google word2vec model. *Jour of Geo Inf Sci* 31(4):825–848
- Yuan NJ, Zheng Y, Xie X, Wang Y, Zheng K, Xiong H (2015) Discovering urban functional zones using latent activity trajectories. *IEEE Trans on Knowledge and Data Engineering* 27(3):712–725

A Evaluated Areas

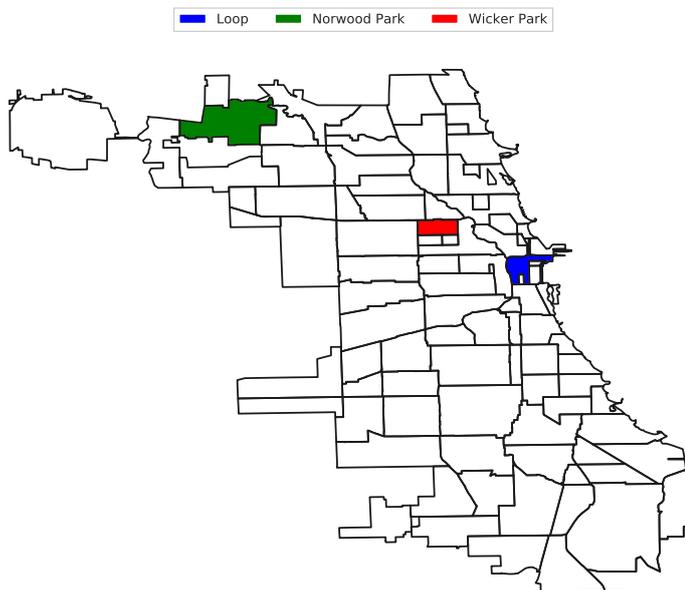


Fig. 21 Evaluated neighborhoods in Chicago, IL.

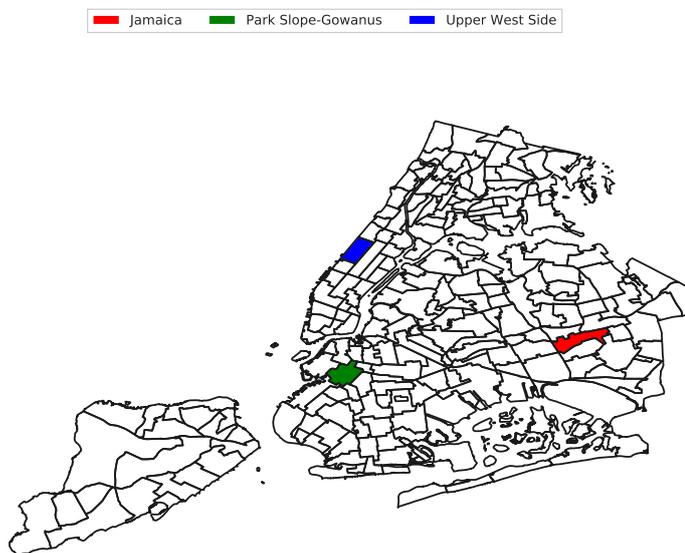


Fig. 22 Evaluated neighborhoods in NYC, NY.

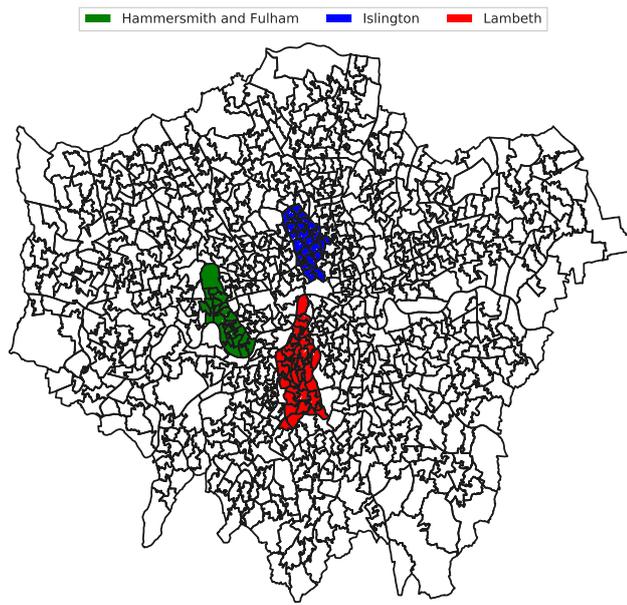


Fig. 23 Evaluated neighborhoods in London, UK.

B Urban Perception Maps

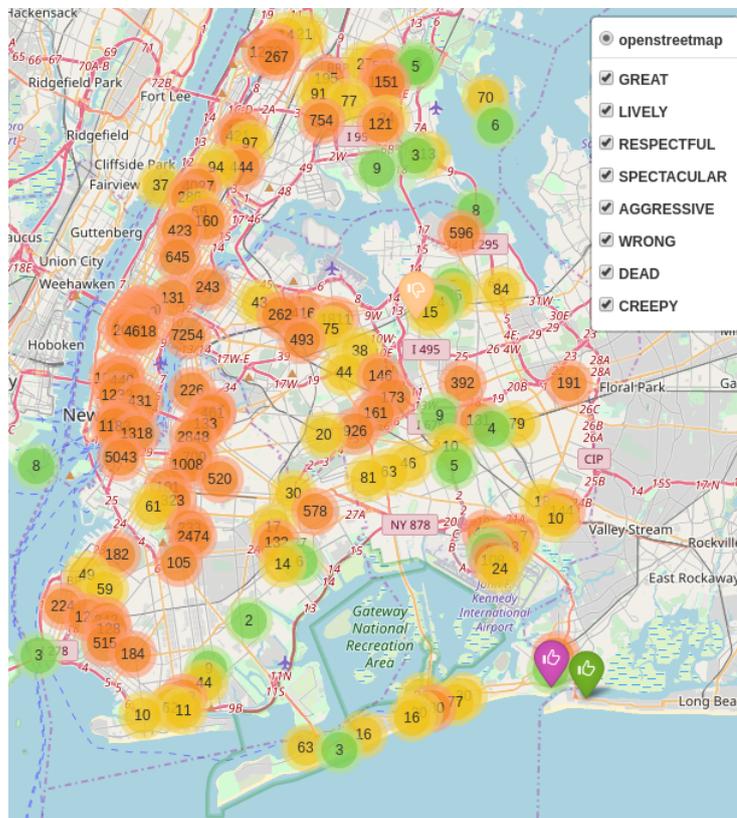


Fig. 24 Screenshot of an interactive urban perception map in New York City. .

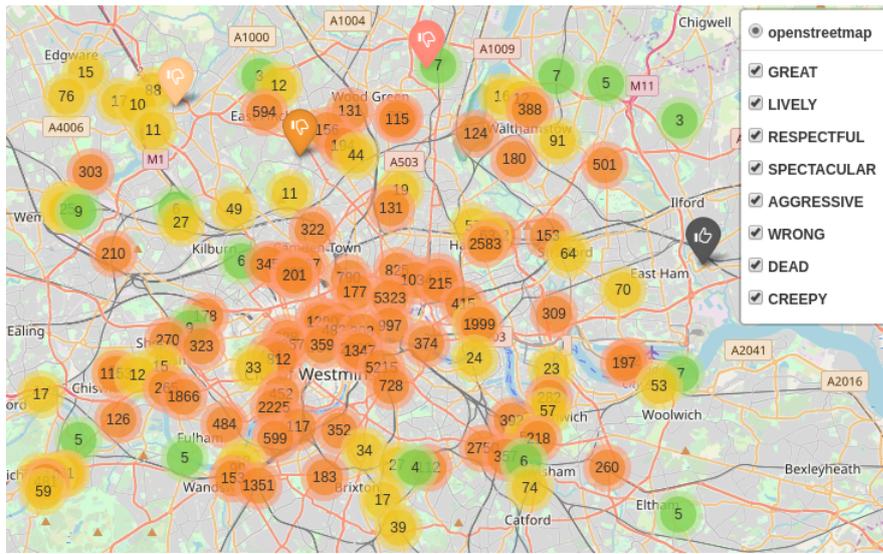


Fig. 25 Screenshot of an interactive urban perception map in London.