GI_Forum 2016, Vol.1
Page: 191-203
Full Paper
Corresponding Author:
scvetojevic@ufl.edu
DOI: 10.1553/giscience2016_01_s191

Positional Accuracy of Twitter and Instagram Images in Urban Environments

Sreten Cvetojevic, Levente Juhasz and Hartwig Hochmair
University of Florida, Fort Lauderdale, U.S.A

Abstract

Twitter and Instagram are social networking services that allow users to share images. To some extent, both platforms provide means for the user to annotate images with geographic location information. Using a selection of images shared through these two platforms, this study compares the photographer's position, which is manually estimated from the scene in the image, with the annotated location information associated with the image and the position of the object being photographed. This approach provides an initial insight into the Twitter user's movement between the location where a picture is taken and the place from where it is uploaded to Twitter. Furthermore, the distance between the photographer's position and the location of the object shown in a Twitter or Instagram photograph can be used to assess the visual prominence of a photographed urban object in relation to its surroundings. Finally, the dataset generated in the research allows us to assess the positional accuracy of location labels in Instagram through comparison of the label position and the true position of the referenced object. For each of the different analyses, this paper discusses sources that could potentially lead to positional errors of images in Twitter and Instagram, and provides a comprehensive set of illustrative examples from different cities.

Keywords:

volunteered geographic information, social media image, positional accuracy, data quality

1 Introduction

Driven by the rapid development in computer, sensor and communication technology, the last decade has experienced a surge in new Web 2.0 and social media applications that allow users to share spatial information over the Web and mobile communication platforms. Two prominent examples of social networking/photo-sharing platforms are Twitter and Instagram. Twitter is an online microblogging service that allows users to send and read short 140-character messages ('tweets'). The nature of Twitter data has been analysed in numerous ways, ranging from the extraction of travel patterns (Hawelka et al., 2014), estimating the influence of socio-economic factors on Twitter activity (L. Li, Goodchild & Xu, 2013), to the "localness" of tweets and other geotagged social media (Johnson et al., 2016). Twitter is also a rich source of images, since users can share links to media from other websites (e.g. YouTube, Instagram), or attach pictures to their posts on Twitter. The spatial aspect of Twitter image-sharing has, however, not so far been discussed in the research literature, although some studies have analysed various other aspects of Twitter images. For example, Thelwall et al. (2015) conducted a content analysis of 800 images tweeted from the UK and the USA, finding that most were photographs, that about 9% of the images displayed mainly text, and that about 15% of images were screen grabs. The same study estimated that about two-thirds of the images were taken immediately before being tweeted. Yanai & Kawano (2014) developed a classifier for grouping streamed Twitter photo data showing food into 100 kinds of food. The results are visualized in a map of prevailing foods, showing popular foods in different parts of Japan. The paper also analysed how the popularity of different dishes, such as "ramen noodle", "curry" and "okonomiyaki", varies by season and region.

The study presented in this paper complements earlier research by assessing the positional accuracy of Twitter images at the urban level. The photographer's position is estimated from the features in the image through visual identification of the location by human analysts. This is then compared to the coordinates of the associated geotagged tweet and the photographed object itself. The visual estimation of the photographer's position from the subject of an image for accuracy assessment of crowd-sourced data has already been applied to data from other photo-sharing services, such as Flickr and Panoramio (Zielstra & Hochmair, 2013). Automated methods to extract the photographer's position from image content have already been developed for regions with high photo density where images overlap sufficiently, and for which a set of control points with known coordinates is provided (Y. Li, Snavely & Huttenlocher, 2010).

Instagram is a photo- and video-sharing platform which allows users to take pictures and videos and to share them with their followers on the Instagram website, as well as through a variety of social networking platforms such as Facebook, Twitter and Flickr. Users can also geotag their shared content. The content and spatial distribution of Instagram images has been analysed in several recent studies. For example, Bakhshi et al. (2014) found that Instagram photos with faces are 38% more likely to receive likes and 32% more likely to receive comments than those without faces. Hochman & Manovich (2013) compared the visual signatures of 13 different global cities using 2.3 million Instagram photos from these cities and used spatio-temporal visualizations of over 200,000 Instagram photos uploaded in Tel Aviv, Israel, to demonstrate how they can offer social, cultural and political insights about people's activities in particular locations and at certain time periods.

Although social media images provide valuable information about a place, the research literature has so far barely touched upon the spatial-accuracy aspect of images shared through Twitter and Instagram. Therefore this paper addresses the following three related research objectives:

- R1: For Twitter images, determine the distance between a photographer's position (derived from the image content) and the geotagged position from which the tweet has been sent. This analysis provides information about a photographer's movement that occurs between taking a picture and sending the tweet with the picture.

- R2: For Twitter and Instagram images, determine the distance between the photographer's position and the photographed object. The range of distances associated with a photographed object give indications regarding the visual prominence of the object.
- R3: For Instagram images, determine the distance between the photographed object and the Instagram location associated with that photograph. This provides information about the positional accuracy of location tags available in Instagram for annotating images with positional information.

2 Study setup

Data collection

This study is based on the local knowledge of human analysts so that the photographer's positions can be estimated from the content shown in Twitter and Instagram images. The study was therefore conducted in geographic areas that participating students, as well as the authors, were familiar with. Since urban environments with their multitude of unique objects (e.g. monuments, stadiums, plazas, churches …) provide more visual clues to estimate a photographer's position than a rural landscape with fewer unique identifiable objects, the study was conducted primarily in urban areas. In addition to the photographer's estimated position, research objective R1 requires the geographic coordinates of the location from which the tweet with an image was sent, and R3 requires the coordinates of the location tag associated with the image by an Instagram user.

Geotagging in Twitter and Instagram

The Twitter mobile application interface allows the user to opt for attaching exact geographic coordinates as metadata along with the tweet. In this case, the geographic coordinates are obtained through the smartphone's geolocation method, which can be based on the built-in GPS receiver, nearby Wi-Fi networks, or the mobile network itself through base station information. The accuracy of the latter method depends on the mobile network infrastructure. As an alternative method of geotagging tweets, the user can pick from a list of nearby locations in the mobile application, where more general geographic entities, such as country, province or city, appear at the top of the list. How general Twitter's suggestions are depends on the geographic region. For example, for photos from Belgrade, Republic of Serbia, the top-most suggested place tag was "Republic of Serbia", whereas for photos from Vienna, Austria, the suggested place tag was "Vienna, Austria". Since the spatial granularity of these places is too coarse for the research tasks proposed in this study, only photos from tweets with geographic coordinates (derived from the mobile phone) were used.

A geotagged image in Instagram does not provide exact geographic coordinates of the location from which the picture was taken, or from which it was sent or uploaded. Instead, it provides the name of the location, selected by the user from a pre-defined list when uploading the image to Instagram. If the photo to be uploaded to Instagram has geographic coordinates in its Exif (Exchangeable image file format) image-file metadata tags, the Instagram application lists locations that are near the coordinates in the Exif metadata. Exif tags contain coordinates if the smartphone geolocation was activated when the image was taken. If the Exif tags do not contain geographic coordinates, the Instagram application lists locations near the current upload location identified by the smartphone. The link to an Instagram image can also be tweeted from within the Instagram application. If the image file that is to be shared via Instagram does not contain geographic coordinates in its Exif metadata and if the smartphone geolocation function is turned off, the image cannot be geotagged. Until a recent change in the Instagram application, users were able to add custom places based on the Exif metadata coordinates or the smartphone position to the list of already-available location names nearby. Therefore a single real-world place, such as a city, state or mountain, can have different Instagram place labels assigned to it, with the same or different coordinates. It is also possible that the same real-world feature is associated with several identical Instagram place labels, where these place labels vary in position. Adding custom place labels in Instagram was deactivated in August 2015.

Obtaining the photographer's position

To obtain the position of the photographer at the time when the picture was taken, we relied on the local knowledge of 47 graduate students who took on this task, for a partial course credit, as part of a GIS graduate course at the University of Florida. Each student was asked to provide us with the bounding boxes ("polygons") of two urban areas they were familiar with, anywhere in the world. For these areas, we collected three types of photos:

1) Photos attached to tweets (hosted by Twitter); links to jpg files are provided in tweet JSON files that can be harvested from the Twitter streaming API.
2) Photos from Instagram shared in tweets (as links to Instagram photos); a tweet contains the link to the Instagram website for that photo. The HTML code of the Instagram website was then parsed for the URL of the corresponding jpg file.
3) Instagram photos: original photos posted on Instagram containing metadata such as user and location information, links to photos, or captions.

Each photo used in the analysis contained at least one type of location information in its metadata. Photos obtained through tweets had geographic coordinates of the place from which the tweet was uploaded. Instagram photos contained a user-assigned location tag. And Instagram photos shared in tweets contained the location of the Instagram location with which users had chosen to annotate it. Instagram images that were either obtained from the Instagram API or sent as a link in a tweet were analysed as one dataset, since for both the only geotagged information available for the image is the Instagram location assigned by the user.

For the data-collection process, in order to obtain a sufficient number of suitable photographs that students could analyse in their selected regions, the areas of the polygons were increased if necessary. This was often necessary for photos attached to tweets (source 1), which occurred in about 7.5% of tweets tagged with exact geographic coordinates. A smaller percentage of geotagged tweets (2.4%) were found to contain links to Instagram images (source 2). The highest photo density in an area was generally obtained from the Instagram API with original Instagram photos (source 3). Prior to distributing photos to students to identify the photographer's position, we manually removed photos that contained profanities and vulgar content. In a Web application that was set up for this study, students could then browse through the photos for their selected urban areas. The task of the assignment for students was to indicate for each image (whenever this was possible) the estimated position of the photographer based on the image content, through adding markers to a “Google My Maps®” map, together with the photo ID. Students were asked to complete this step for 20 images from each data source. If this was not possible, they were asked to analyse more images from any data source (whichever one worked) to reach a total of 60 images. The marker locations indicated by students were then extracted from the shared “Google My Maps®” maps using a customized Python script and inserted into a PostgreSQL database. The authors of this paper went through the same steps for selected areas in Vienna, Salzburg, Budapest, Szeged, Ispra and Belgrade. For the next steps, the photos from only 23 students (out of the original 47) were further processed and analysed to reduce the time-consuming process of data cleaning.

That is, for quality assurance all of the photographer positions indicated by the 23 students were manually checked by the authors in a customized Web application that showed the original photo content, the specified position on a map as a marker, and the “Google Street View®” image for that position next to the map (where available). The Web application enabled us to accept the photographer’s position as indicated by the student, to move the marker position, or to exclude a photo if it was obviously misplaced and if we could not identify the photographer’s correct position based on the satellite image view or “Google Street View®”. Based on these data it was possible to measure the distance between the photographer’s position and (a) the geotagged position of the tweet containing the picture, and (b) the location position associated with an Instagram photo. In addition, for photographs that showed a focus on an object that could be well approximated through a point location, such as a clearly discernible building, the authors placed markers at these locations. If it was unclear which object the photographer was focusing on, the photo was excluded.

Table 1 summarizes the number of photographer positions obtained per country and source that were retained for further analysis. Values in parentheses indicate the number of object locations that were identified by the authors. Depending on the specific research objective, different data columns are used from Table 1 (see section 2.2).

Table 1: Number of identified photographer positions and object locations (in parentheses).

Country	Twitter	Twitter/Instagram	Instagram	Total
Austria	45 (24)	50 (27)	54 (16)	149 (67)
Canada	26 (6)	28 (1)	26	80 (7)
Germany	1	2	18	21
Haiti	1		16 (4)	17 (4)
Hungary	11 (5)	40 (14)	68 (25)	119 (44)
India	7 (2)	12 (3)	14 (2)	33 (7)
Italy		2	41 (6)	43 (6)
Kenya	3	3	3 (1)	9 (1)
Libya	3 (1)	3 (1)	52 (11)	58 (13)
Puerto Rico	1	13 (3)	4 (1)	18 (4)
Serbia	24 (12)	19 (8)	26 (13)	69 (33)
Slovakia	5 (1)	16	18	39 (1)
Turkey	10	12 (1)	31 (10)	53 (11)
United Arab Emirates	4 (1)	8 (2)		12 (3)
United Kingdom	6 (1)	18 (5)	16 (6)	40 (12)
United States	126 (21)	203 (32)	546 (59)	875 (112)
Total	**273 (74)**	**429 (97)**	**933 (154)**	**1635 (325)**

Figure 1 plots the photo locations from Table 1; Figures 2 and 3 provide a zoomed view of available data sources for parts of Vienna and Belgrade.

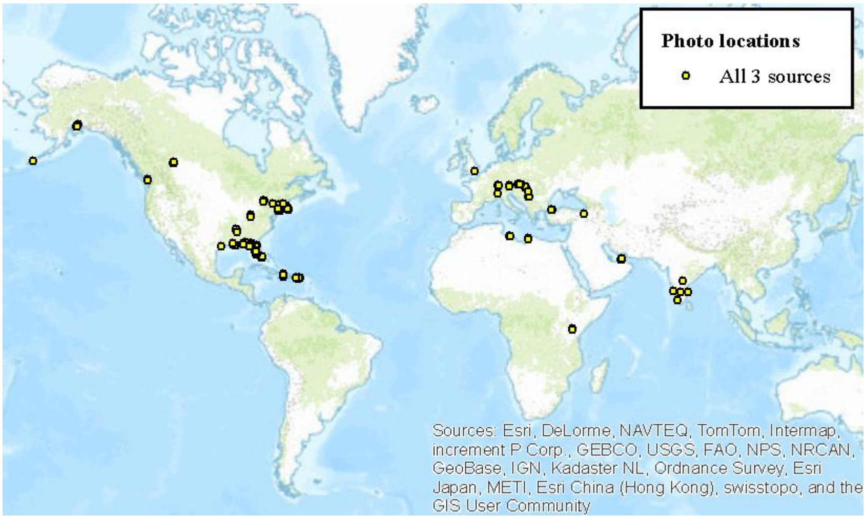

Figure 1: General locations of photographs analysed

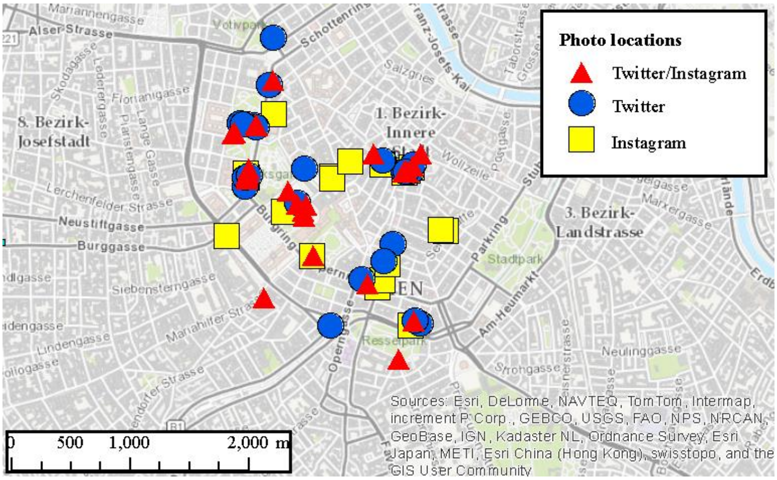

Figure 2: Twitter and Instagram photo positions in Vienna

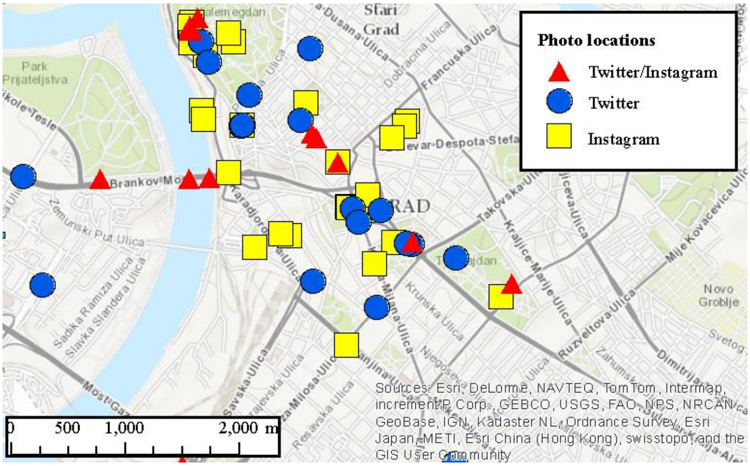

Figure 3: Twitter and Instagram photo positions in Belgrade

Data analysis

The analysis consists of three parts, corresponding to the three research objectives. We estimated the movements of Twitter users between taking a photograph and uploading it to the Twitter site (R1) by measuring the distance between these two positions. To assess regional differences, each data point was assigned to a geographic area, i.e. North America (including the Caribbean), Europe, and other. The dataset consists of 273 individual features from Twitter images.

To answer R2, which assesses the visual prominence of objects, a dataset containing 325 Twitter and Instagram photos was used, for which the positions of both the photographer and the photographed objects could be identified. We hypothesized that both its type and its visual context affect the visual prominence of an object. Therefore each photograph was assigned to one of the following categories: (a) prominent building, spatially separated from other buildings; (b) object photographed from a location separated from the object by water, e.g. through a fountain or across a river; (c) all other photos. The last group included, for example, pictures of local businesses in city-centre locations, or other points of interest such as small monuments or fountains.

For R3, which analyses the accuracy of the Instagram location by measuring the distance between the photographed object and the annotated Instagram location, the dataset comprised 251 photos. This dataset is a subset of the dataset used to answer R2, containing only photos originating from the Instagram platform.

3 Analysis results

R1: Positional accuracy of Twitter images

The Twitter dataset can be used to study the movement of a photographer between taking a picture and uploading it to Twitter. The distribution of distances between photo and upload locations follows approximately a power law function with an exponent value of 0.80 and an R-Squared of 0.66 when binning distances into 1 km categories (Figure 4a). The log-log plot reveals that more than 60% of photos were uploaded within a 1 km radius of the original photo location. At the other end of the range, 2% of all photos were uploaded more than 100 km away from the place where they were taken.

Different user patterns could be observed for posting photos on Twitter. Approximately 30% of the photos were posted within 50 m of the actual location. This distance closely resembles the maximum smartphone positioning error in urban environments, and so these photos can be considered as instant uploads. 10% of photos were posted from more than 10 km away from the original location. This category includes, for example, holiday images or photos from sporting events held in different cities. Users in this category did not upload their photos instantly. The spatial distribution of the intermediate-distance category provides some information about the locations from where social media users post their photos. In some cases, when the offset is great, the upload position corresponds to possible open Wi-Fi hotspots and hotels. This might be indicative of tourist Twitter activities, for example when tourists do not have a mobile phone plan that includes uploading data while abroad and are therefore unable to upload their photos instantly. Images are often uploaded from areas that appear to be residential, but have been taken somewhere else, e.g. in city-centre locations.

Since distances in the three global regions compared do not follow a normal distribution, even after using a log transformation, a non-parametric test was applied to test the effect of geographic region on median distance offsets. Data points were categorized into North America/Caribbean (AME), Europe (EUR) and other (OTH, consisting of locations from Arabic countries, India and Kenya). Descriptive statistics of distances for these categories can be found in Table 2. These reveal that median distances, which are not affected as much by outliers caused by tweets from other cities as the mean distance, are highest for regions outside North-America/Caribbean and Europe. Results of the Mood's median test show that the geographic region has a significant effect on the distance between the photo and upload locations ($p = 0.02$). This can, potentially, be explained by differences in Wi-Fi and mobile data infrastructure, which has generally better coverage in regions of stronger economic development, requiring users in less developed countries to move further for internet connection and sending a tweet. Figure 4b shows boxplots of the log-transformed data grouped by geographic region, supporting the pattern from Table 2.

Table 2: Descriptive statistics of distances between photo upload and photo position in different geographic regions

Region	Mean [m]	Median [m]	SD [m]	N
North America and the Caribbean	7389.0	198.7	20606.4	154
Europe	2837.0	627.7	13077.5	92
Other	3668.0	1559.0	6983.1	27

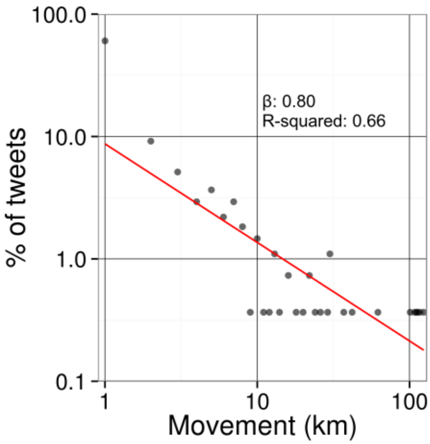

(a)

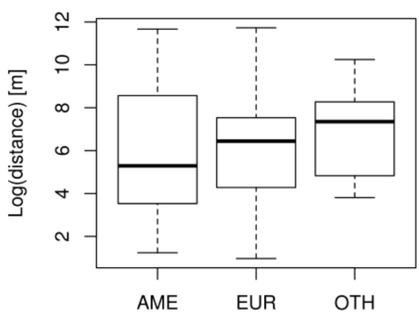

(b)

Figure 4: Power law function fitted to the distance between upload and original photo location in Twitter (a), and boxplots of distances in different geographic regions (b)

R2: Distance between photographer and object

The distance between the photographed object and the photographer can be interpreted as the visual prominence of an object, with larger values indicating that the object can be seen (and is interesting enough to be photographed) from further away. Only photos that have a clear focus on an object were used. Therefore landscapes, city panoramas, portraits and photos with scenery were excluded from the analysis. Visual inspection of the distance data revealed that most photos were taken in close proximity to the photographed object, because urban environments usually prevent distant views due to the high building density. Figure 5a shows a typical image setup in a city, with many objects, such as stairways (lower left inset), being photographed in close-up. The figure also shows that photos of landmark buildings

tend to be taken from greater distances, which is because of their visual prominence and their contexts, which often include large squares or parks. A similar case occurs if a water body is located between the object and the photographer (Figure 5b), preventing the user from moving closer, and often providing a scenic foreground for the photograph. Boxplots of distances for these categories are shown in Figure 5c. A one-way ANOVA test on the log-transformed distances indicates a significant effect of the object category on the photographer's distance ($F(2,322) = 87.47$, $p < 0.001$).
The overall distribution of distances between object and photographer also follows a power law function with an exponent value of 1.31 and R-Squared of 0.89 (Figure 5d). Of the 325 social-media photos analysed, 47% of those with identified objects were taken within 25 m of the object, and only 16 % were taken more than 100 m away from the object.

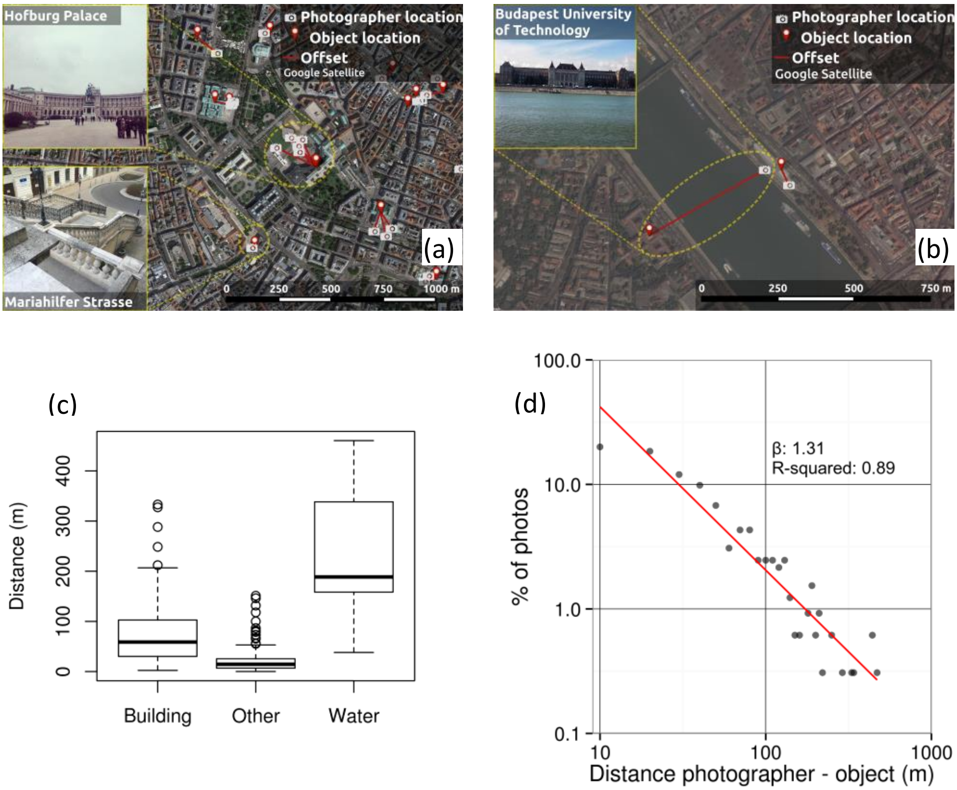

Figure 5: Offset between the photographer and identified object in Vienna (a), and in Budapest (b); boxplot of distances for different object categories (c); power law function fitted to the frequency distribution of distances for Twitter and Instagram photos (d)

R3: Distance between Instagram location and object position

Instagram location labels are diverse in nature and can denote, among other things, physical objects, such as a building, street or monument, or administrative units, such as a city. Users previously had the ability to create custom locations, which resulted in a high density of distinct Instagram locations in urban environments, as shown in an example for Salzburg (Figure 6a). Since August 2015, however, attaching photos to existing locations is now the only way to geocode Instagram photos. The updated app also prevents users from creating new locations inside the Instagram app itself. The offset between the identified objects and the Instagram locations in the dataset analysed ranges from 2 m to 24 km (median: 85 m, mean: 635 m). 52 % of the locations were less than 100 m from the object and 14 % of them were more than 1km away from it.

There are several possible reasons for location offset errors. Among the locations more than 1 km away from the identified object, several were tagged with general names, such as the name of a town (e.g. Ispra) or a geographic area (e.g. Dutch Harbor). This is not necessarily a positional error of the Instagram location, but may reflect, rather, the user's inclination towards increased privacy (i.e., obscuring his or her exact location), lack of local knowledge, the thinking that a general location name is the best fit for describing the photo content, or the absence of an appropriate Instagram location nearby. Another explanation for large offsets that are not related to Instagram location position errors is that the user has mistakenly picked the wrong location label for the photo. If the photo is not tagged with geographic coordinates in its Exif tags, users rely on the Instagram location suggestions that are based on their current position. In such cases, when a user moves away from the photo location, an image may become associated with a place in the proximity of the upload location, which can be some distance away. Figure 6b shows an extreme example of a photo (distance between location given and actual object: 3.3 km) that is associated neither with the place from where it was taken, nor the true location of the object shown on the photo, but with a third location, which is most probably close to the place of upload (the northern most point).

Furthermore, a number of Instagram photographs with a large distance between the true position of the object photographed and the position of the Instagram label revealed misplaced Instagram labels, where locations do not correspond with their true positions. An example is provided in Figure 6c, where Instagram locations are marked as red dots. In these cases, it is possible that the user who created the location travelled a considerable distance south-east before creating a custom location. This phenomenon implies that custom locations were geotagged based on the smartphone's geolocation, i.e. the current position of the user. This is illustrated in an example for St. George Island, Florida (Figure 6d). The spread of Instagram locations around the true position of the object (a lighthouse) implies that locations were most likely added by Instagram users, with coordinates corresponding to their smartphone locations. The example also shows that the same object can have multiple Instagram locations. One problem with misplaced locations is that users can add photos to them without being aware of the position error, since only the location names are shown in the apps, not their map location.

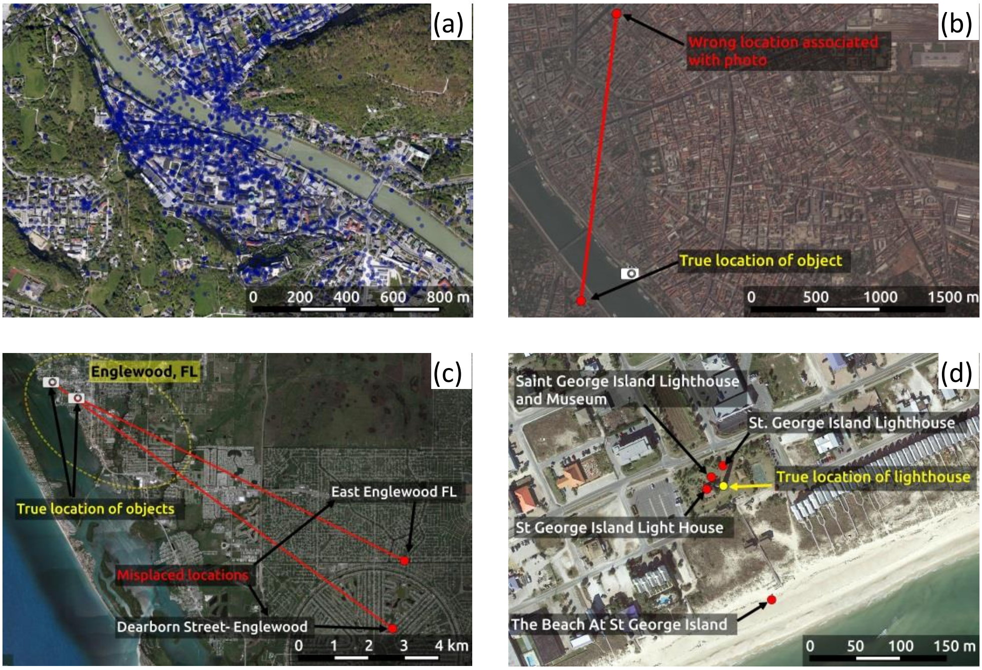

Figure 6: Spatial distribution of Instagram locations in Salzburg (a), incorrect selection of Instagram location in Budapest (b), misplaced Instagram locations in Florida (c), multiple locations for the same object with similar labels in St. George Island, Florida (d)

4 Conclusions and future work

This study analysed the positional accuracy of geotagged images shared over Twitter and Instagram, using the photographers' estimated positions from the image content. For Twitter, the analysis provided some explanations for observed patterns of distance offsets between photo capture location and tweet position, including Wi-Fi availability. The study considered primarily images taken within urban areas, in order to facilitate recognition by the analyst. Offset distances between photo capture location and tweet position can be expected to be much larger if distances to locations outside the city limits (e.g. even in other countries) are taken into account as well. Extending this kind of analysis to the global scale is part of the plans for future work. The study showed that Twitter and Instagram images help to identify the visual prominence of objects, which is affected by the type of object and its immediate surroundings. The analysis therefore relates to the purely visual aspect of a landmark's attractiveness; analysis could be expanded to determining the semantic and structural attraction of landmarks (Raubal & Winter, 2002). The study also provided various explanations for inaccuracies observed in Instagram location labels, such as travel between the location where a picture was taken and the location where it was uploaded. For future work, we plan to explore the density and accuracy of place labels in greater depth for cities around the world, and to relate their spatial characteristics to those of other place label collections, for example in Foursquare/Swarm.

References

Bakhshi, S., Shamma, D. A., & Gilbert, E. (2014). Faces Engage Us: Photos with Faces Attract More Likes and Comments on Instagram CHI '14 Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. New York, N: ACM.

Hawelka, B., Sitkoa, I., Beinat, E., Sobolevsky, S., Kazakopoulos, P., & Ratti, C. (2014). Geo-located Twitter as proxy for global mobility patterns. Cartography and Geographic Information Science, 41(3), 260-271.

Hochman, N., & Manovich, L. (2013). Zooming into an Instagram City: Reading the local through social media. First Monday, 18(7). doi: 10.5210/fm.v18i7.4711

Johnson, I. L., Sengupta, S., Schöning, J., & Hecht, B. (2016, March 07 - 10). The Geography and Importance of Localness in Geotagged Social Media. Paper presented at the ACM International Conference on Intelligent User Interfaces, Sonoma, CA.

Li, L., Goodchild, M. F., & Xu, B. (2013). Spatial, temporal, and socioeconomic patterns in the use of Twitter and Flickr. Cartography and Geographic Information Science, 40(2), 61-77.

Li, Y., Snavely, N., & Huttenlocher, D. P. (2010). Location Recognition Using Prioritized Feature Matching. In K. Daniilidis, P. Maragos & N. Paragios (Eds.), Computer Vision – ECCV 2010 (Vol. LNCS 6312, pp. 791-804). Berlin: Springer.

Raubal, M., & Winter, S. (2002). Enriching Wayfinding Instructions with local landmarks. In M. Egenhofer & D. Mark (Eds.), GIScience 2002 (Vol. LNCS 2478, pp. 243-259). Boulder, Colorado, USA: Springer.

Thelwall, M., Goriunova, O., Vis, F., Faulkner, S., Burns, A., Aulich, J., . . . D'Orazio, F. (2015). Chatting through pictures? A classification of images tweeted in one week in the UK and USA. Journal of the Association for Information Science and Technology. doi: 10.1002/asi.23620

Yanai, K., & Kawano, Y. (2014). Twitter Food Photo Mining and Analysis for One Hundred Kinds of Foods. In W. T. Ooi, C. G. M. Snoek, H. K. Tan, C.-K. Ho, B. Huet & C.-W. Ngo (Eds.), Advances in Multimedia Information Processing – PCM 2014 (Vol. LNCS 8879). Berlin: Springer.

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.