

Using Reverse Viewshed Analysis to Assess the Location Correctness of Visually Generated VGI

Hansi Senaratne,* Arne Bröring[†] and Tobias Schreck*

*University of Konstanz

[†]Institute for Geoinformatics, University of Münster & 52°North Initiative for Geospatial Open Source Software GmbH

Abstract

With the increased availability of user generated data, assessing the quality and credibility of such data becomes important. In this article, we propose to assess the location correctness of visually generated Volunteered Geographic Information (VGI) as a quality reference measure. The location correctness is determined by checking the visibility of the point of interest from the position of the visually generated VGI (observer point); as an example we utilize Flickr photographs. Therefore we first collect all Flickr photographs that conform to a certain point of interest through their textual labelling. Then we conduct a reverse viewshed analysis for the point of interest to determine if it lies within the area of visibility of the observer points. If the point of interest lies outside the visibility of a given observer point, the respective geotagged image is considered to be incorrectly geotagged. In this way, we analyze sample datasets of photographs and make observations regarding the dependency of certain user/photo metadata and (in)correct geotags and labels. In future the dependency relationship between the location correctness and user/photo metadata can be used to automatically infer user credibility. In other words, attributes such as profile completeness, together with location correctness, can serve as a weighted score to assess credibility.

1 Introduction

In today's information-driven society, Volunteered Geographic Information (VGI) has increased immensely over the past years. With this massive increase of data production by volunteers, the need for caution about data credibility becomes ever more pressing. Humans perceive and express geographic regions and spatial relations imprecisely, and in terms of vague concepts (Montello et al. 2003). This vagueness in human conceptualization of location is due not only to the fact that geographic entities are continuous in nature, but also to the quality and limitations of spatial knowledge (Hollenstein and Purves 2011).

Hovland et al. (1953) expressed *credibility* as the believability of a source or message, which comprises primarily two dimensions, the *trustworthiness* and *expertise*. Flanagan and Metzger (2008) further asserted that, while trust and expertise have different meaning from credibility as well as from each other, one conceives of credibility as possessing a combination of both trust and expertise. Hence, due to the subjective and objective nature of trust and expertise, credibility is a complex concept that has to do with the *believability* of a source. Therefore, in assessing the credibility of data, one needs to consider factors other than accu-

Address for correspondence: Hansi Senaratne, Department of Computer and Information Science, University of Konstanz, Box 78, 78457 Konstanz, Germany. E-mail: hansi.senaratne@uni-konstanz.de

Acknowledgements: The research leading to these results was supported by DFG Research Training Group GK-1042 "Explorative Analysis and Visualization of Large Information Spaces", University of Konstanz. Further, we're thankful for the reviewers' comments that helped in improving this article.




Figure 1 Geotags of Flickr photos that were textually tagged as “Angkor” and “Cambodia”

racy that contribute to this perception of trustworthiness and believability. Metadata about the origin of VGI can provide a foundation for judgment on the quality and trustworthiness (Frew 2007).

In the case of Flickr, as an example of a platform for visually generated VGI, volunteers can upload photographs to share with others. A Flickr user can maintain a profile to which uploaded photos are linked and state metadata such as his/her real name, the date of registration, hometown, or contacts to other users. Also, metadata for the picture itself can be specified, such as title, caption, textual tags describing the photo (label), or the dates of capture and upload. Additionally, a spatial reference of the photo can be given in the form of geographic coordinates. This *geotag* can either be produced by an external GPS device, automatically recorded with a camera built-in GPS, or it can be manually located using Flickr’s map interface at varying levels of resolution (i.e. neighborhood, city, country).

As well as the geotag, which consists of geographic coordinates, Flickr users often specify the place of interest to which the picture relates, as a textual tag. The map shown in Figure 1 displays all geotags of Flickr photos annotated with the textual tags “Angkor” and “Cambodia”. Although most of the photos of this data set are geotagged within the area of this ancient city in Cambodia, the visual analysis shows that there are also many pictures located far away from it. For example, one photo displaying a part of Angkor is geotagged at a location in California: <http://www.flickr.com/photos/rbleib/5030263322/in/set-72157624911484519/>. Becker and Bizer (2011) also demonstrated, through their work on the *Flickr Wrappr*, how pictures on Flickr are incorrectly geotagged.

In this article, we describe a concept for assessing the location correctness of visual VGI content based on a reverse viewshed analysis. The basic idea entails validating the location of a described object within a user-provided image by testing whether that object can be viewed from the position where the photo is geotagged. We propose this approach to validate correctness of geotagged photographs provided by Flickr as a VGI data source. This approach generalizes to other visual VGI data sources as well.

We test our concept by experimental analysis. In our experimental setup, we downloaded metadata of photographs for two points of interest (POI), which are textually tagged as “Brandenburg Gate”, “Berlin” and “Reichstag”, “Berlin”. Further, in order to derive a reference quality measure for location correctness, a reverse viewshed analysis for these POI is calculated. A reverse viewshed analysis holds the same principles as the viewshed analysis, however, it is utilized to determine the visibility to a given POI from many observer points

(Fisher 1996). In a third step, we were able to determine which photographs are textually tagged with the description of the given POI (e.g. “Brandenburg Gate” and “Berlin”) *and* which are correctly geotagged within the range of visibility to that point. If the POI does not fall within the area of visibility from the geotagged image, then either the image misrepresents the location from where the photograph was taken, or the photographed content represents something else other than the POI but is tagged as the latter. Photographs belonging to either of these two groups are considered to be tagged with incorrect locations.

Using this approach, we investigate which metadata of photographs (e.g. tag count of photographs, comment count of photographs, etc.) as well as metadata about users (e.g. the number of photos, the number of contacts, or the camera used) can be utilized to eventually infer the credibility of photographers regarding a correct geotagging. We achieve this through analyzing the relationship between those metadata and the location correctness as the reference quality measurement. For the future, we have paved the road with this approach for new applications that can automatically assess quality of Flickr photographs. This methodology can also be transferred to other visual VGI sources (e.g. Panoramio).

The remainder of this article is organized as follows. In Section 2 we review selected work related to our research. In Section 3 we discuss our approach in detail. Section 4 analyzes a sample dataset and derives observations on dependencies between user/photos metadata and location correctness. We then discuss results and limitations in Section 5 and conclude with further ideas for future work in Section 6.

2 Background and Related Work

With the Web 2.0 in place, citizens can contribute data in the form of text, audio, or video on the Web, making the consumers of data also the producers. This is termed as User Generated Content (UGC). Surowiecki (2005) shows how a group of people may contribute to a solution of a problem that an expert may be unable to solve. A special case of UGC is where citizens, quite often untrained, create *geographic* information which may or may not be accurate, on dedicated web platforms, e.g. OpenStreetMap (<http://www.openstreetmap.org>), Wikimapia (<http://www.openstreetmap.org>), Google MyMaps (<https://www.google.com/maps/mm>), and Flickr (<http://www.flickr.com>). Goodchild (2007) coined this phenomenon as Volunteered Geographic Information (VGI). As of January 2012, Flickr has reported hosting over six billion images (<http://www.searchenginejournal.com/the-growth-of-social-media-an-infographic/32788/>). Around 3% of the Flickr images were geotagged in 2009 (<http://code.flickr.net/2009/02/04/100000000-geotagged-photos-plus/>), and OpenStreetMap statistics (http://www.openstreetmap.org/stats/data_stats.html) state that over a million registered users have contributed more than three billion track points around the world. Rinner et al. (2008) identified an exponential growth for such VGI.

When consuming VGI, it is important to keep in mind that the content is not quantified by the objective notions of data quality, nor does it rely on traditional authorities who enforce data quality standards (Flanagin and Metzger 2008). Instead, the *credibility* of the data depends on the personal accuracy of the producers.

2.1 Credibility of VGI

Extensive research has been conducted to assess the credibility of user generated geo information on different platforms. The central question has been whether we can trust the data volunteers to

produce data of *usable* quality for convenient usage and to derive accurate conclusions. A few VGI platforms have taken measures to moderate the credibility of user-generated data to ensure reliability. For example, the Audubon Society's Christmas Bird Count (<http://birds.audubon.org/christmas-bird-count>) is open for bird watchers to observe migration patterns, bird population, etc., and to contribute these observations on the open platform. Before they can take part, the users should possess domain knowledge to a certain degree and are given proper training. Project GLOBE (<http://training.globe.gov/>) is another example; it encourages school children around the world to observe and collect their local weather data and contribute it to a common platform. To ensure data credibility, the supervising teachers are given thorough training on data collection and uploading so that they can guide the students through the process. WikiScanner (<http://wikiscanner.virgil.gr/>) is an example for assessing *user* credibility in UGC. WikiScanner cross references the edits on Wikipedia (<http://www.wikipedia.org>) with the data on the editor of the associated block of IP addresses of various organizations, although Wikiscanner does not distinguish between edits made by authorized users from IP addresses originating from organizations and edits made by unauthorized intruders and users of public access computers. It is author identity that provides credibility judgment.

Extending the work by Haklay (2010), Girres and Touya (2010) assessed the quality of OpenStreetMap data for France by comparing them with officially surveyed data. They assess the quality of these VGI within five GI quality components; *accuracy* (positional, thematic, temporal, semantic), *completeness*, *logical consistency*, *lineage*, and *usage*. In their research on the quality of OpenStreetMap data, Haklay et al. (2010) found that the positional accuracy of features improves as the number of editors increases, up to 13. Goodchild and Li (2012) proposed a three tier approach to assuring VGI quality: (1) crowd-sourcing (number of contribution and accuracy); (2) volunteers who are given roles in the *hierarchy* to moderate the data accuracy; and (3) the geographic approach, where geographic features on a map are inferred from knowledge of the surrounding geography.

Ciepluch et al. (2010) assessed the accuracy of OpenStreetMap data based on completeness of the map, currency of the spatial information, correctness with relation to the ground truth data and local knowledge. The authors assert that in order for OpenStreetMap to be taken seriously, quantifiable metric measurements must be evaluated for the OpenStreetMap accuracy and coverage. Furthermore, Bishr and Kuhn (2007) state that the lack of quality measures can affect the usability of user contributed data, and that trusted users provide more useful data. This issue led Goodchild (2009) and Coleman et al. (2009) to categorize the volunteers of VGI into different groups based on their knowledge and experience with geo information.

2.2 Classification of Users to Assess Data Credibility

Goodchild (2009) classified data producers as falling into either *Neo Geography* or *Academic Geography*. Neo Geography is where the role of the user intersects between the roles of subject, producer, presenter and consumer. In other words, the volunteer does not clearly belong to any of these distinct roles. In contributing to VGI, however, they are all experts in their own communities. Volunteers in the academic geography category, on the other hand, are involved in professional geography, e.g. as surveyors or cartographers. Coleman et al. (2009) classified data volunteers as overlapping between *Neophytes*, *Interested Amateur*, *Expert Amateur*, *Expert Professional*, and *Expert Authority*. He analyzed these groups on the basis of what motivates them to produce data on such sources. Coleman et al. (2009) further implied that volunteers fall into the above categories depending on three different contexts: *Market driven*, *Social networks*, and *Civic/Governmental*. Volunteers who fall into the category of

Market driven contribute data on commercial databases or services such as TomTom (<http://www.tomtom.com/>) or Garmin (<http://www.garmin.com/us/>). Volunteers falling into *Social Networks* contribute to sources such as OpenStreetMap, Flickr, etc. Volunteers falling into *Civic/Governmental* contribute data out of concern for their city/society, for example to PPGIS (<http://www.ppgis.net/>).

Zwol (2007) presents a characterization of user behavior on Flickr, and shows that the number of contacts per user and the number of pools an image belongs to can be used to predict the *popularity* of a photo. He further asserted that the social affiliation sustained by the network of contacts within Flickr is important for the popularity of photos. In other Flickr analyzes, Friedland et al. (2011) as well as Moxley et al. (2008) utilize textual tags of Flickr content along with certain visual cues to determine the geographical coordinates of the place being captured in the visual content.

2.3 Tagging Behavior in Flickr

Flickr photos have been explored in a multitude of geographical analyzes. For instance, Jankowski et al. (2010) and Crandall et al. (2009) explored spatial and temporal patterns in user movement and their interests in landmarks and events captured through Flickr. These Flickr photos are organized or searched with the help of accompanying tags that come in various forms. Ames and Naaman (2007) have comprehensively discussed the concept of *tagging* and have identified two main incentives that motivate users to tag: (1) *sociality*, describing who is intended to use the tag; and (2) *function*, describing the intended usage of the tag, which could be either for organizational or retrieval purposes, and also to gain attention for the tagged content. Tagging an image is a means of adding metadata to the content in the form of specific keywords to describe the content (Golder and Huberman, 2006), or in form of geographic coordinates (Geotagging) to identify the location linked to the image content (Valli and Hannai 2010). Moxley et al. (2008) developed a tool that suggest tags for a given image, based on the geographic context and visual relevance. Crandall et al. (2009) analyze the content of a photo based on text labels and image data, and the structure based on the geospatial data. They further assert that within a street level scale, text tags alone can be a useful source for estimating the location, but in combination with visual cues it can be an even stronger component in validating the location. Furthermore, while Girardin et al. (2008) analyzed tags of Flickr photos to explore how people perceive their environment and the underlying semantics on how they describe the urban space, Sigurbjoernsson and Zwol (2008) found in their study of selected Flickr photos that most frequently, tags represent a location followed by artifacts/objects.

Building up on these works, we introduce the assessment of location correctness of geotagged Flickr photographs based on visibility. Specifically, the reverse viewshed analysis is proposed as an objective baseline measure for positional accuracy which can serve for additional investigations into which *characteristics* of a VGI volunteer influence the credibility of his/her contributions. We take Flickr as the experimental data source. Our approach however, is more generic and applicable to estimate positional credibility for any VGI data source where geocoordinates and textual image tags which denote an object or place of interest, occur. Our approach is discussed in the following section.

3 Approach

In this section, we first provide an overview of the proposed approach to evaluating the location correctness of visually generated VGI (Section 3.1). Sections 3.2 and 3.3 then describe our

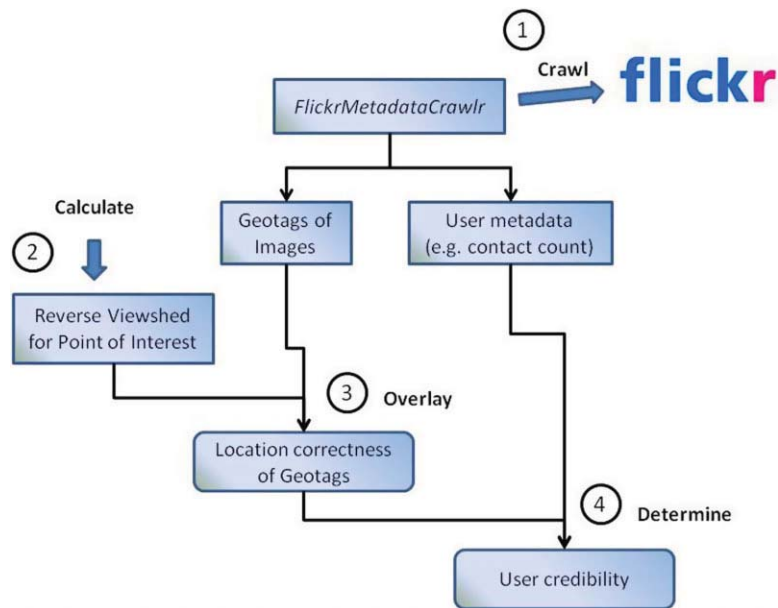


Figure 2 The workflow of the proposed methodology to assess the location correctness of visually generated VGI. We utilize Flickr as the data source

methodology in detail; we describe the computation of a reverse viewshed for geotagged Flickr images for Brandenburg Gate and Reichstag in Berlin followed by the implementation of a crawler to fetch metadata for Flickr geotagged images, through the Flickr API. At the end of the section, we present an overlay between viewsheds and geotags of selected Flickr images to depict the location correctness of the respective geotagged images.

3.1 Overview on Assessing the Location Correctness of Visually Generated VGI

Our approach to assess the location correctness of visual VGI entails a series of steps as illustrated in Figure 2.

In a first step, the metadata of all photographs for our POI, i.e. Brandenburg Gate and Reichstag in Berlin, are automatically downloaded by a crawler. The label of the POI is part of the textual tags for each of those photographs. The fetched metadata includes the latitude and longitude, the geotag of the image. In a second step, a reverse viewshed is calculated for the POI. A reverse viewshed successively determines from which observer points the POI is visible. This allows cross validating if a picture was taken within the vicinity to the POI. In a third step, we assess the location correctness of the geotags based on this visibility analysis. Images belonging to observers whose line of sight did not include the position of the POI are regarded as incorrectly geotagged, and images belonging to observers whose line of sight includes the positions of the POI are regarded as correctly geotagged. In a fourth step, we look into the various user/photo metadata attributes of the photograph to explore how these can be used in association with the reverse viewshed, to automatically classify VGI producers concerning accurate geotagging.

3.2 *The Reverse Viewshed Analysis*

Within our experiment we chose the Brandenburg Gate and the Reichstag of Berlin as the two points of interest. In the following section we compute the reverse viewshed for these POI.

A viewshed analysis can be conducted in a standard Geographic Information System (GIS), to determine the total area that is visible from a given point (O'Sullivan and Unwin 2003). Viewshed analysis is carried out in a variety of applications including but not limited to urban environment planning (Lake et al. 1998), locating telecommunication towers (De Floriani et al. 1994), or tree cover conservation (Sherren et al. 2011). A viewshed of a particular point is calculated from elevation data around the region, which is employed in an algorithm that estimates the difference of elevation of the intermediate pixels between the viewpoint and the target pixels. In order to determine the visibility of the target pixel, the intermediate pixels are analyzed for their line of sight (line of sight determines if the target pixel is visible from the viewpoint, or obscured). If the line of sight is visible then the target pixel is included in the viewshed, if obscured then the target pixel is not included in the viewshed (Kim et al. 2004). Among many who developed efficient viewshed algorithms, (e.g. Fisher 1991, 1993, Wang et al. 1996), Fisher (1996), Kidner et al. (1999), and Ralling et al. (1999) also discussed *reverse* viewshed analyses. A reverse viewshed analysis holds the same principles as the viewshed analysis. However, it is utilized to determine the visibility of a given target point from *many* observer points (Fisher 1996). Fisher (1996) distinguished between the area which can be seen from the location (viewshed) and the area from which a location can be *viewed* (reverse viewshed), based on the height differences between the viewing point and the viewed object. Taking this into consideration, we have utilized the same technique but a different procedure to generate a viewshed; i.e. instead of taking one viewshed from the target point, we create multiple viewsheds from the observer points to validate whether the target falls within the visibility of the observer. We use this reverse viewshed analysis to determine the visibility of the Brandenburg Gate or the Reichstag, respectively, from the surrounding observer points. This is discussed in more detail in the following section.

3.3 *Accessing Metadata of Visual VGI*

To make metadata of visual VGI available to the developed process and the viewshed analysis, we have implemented a tool for the Flickr example, the so-called FlickrMetaCrawlr (the source code of our *FlickrMetaCrawlr* can be accessed at: <http://ifgi.uni-muenster.de/~arneb/FlickrMetaCrawlr.jar>). This tool is able to programmatically download metadata of Flickr photos and users. The FlickrMetaCrawlr therefore relies on the open Flickr API (<http://www.flickr.com/services/api>) and fetches metadata of Flickr photographs for a specified set of tags.

The Flickr API restricts applications to access a maximum of 5,000 photos in a single API query execution. However, a certain tag combination may result in a much larger number of photos – e.g. searching for *Times Square* and *New York* results in around 15,000 geotagged photos. Hence, a mechanism has to be included that divides the initial query into sub-queries which result in less than 5,000 photos. Therefore, to facilitate access to all photographs that confine to a tag query, the FlickrMetaCrawlr utilizes a *quadtree* algorithm (Samet 1984).

The quadtree is essentially applied to the geographic space and subdivides it recursively into four quadrants starting with the maximum extent (the bounding box between 180°W, 90°S and 180°E, 90°N). A division into four quadrants is performed in case more than 5,000 photos are contained within a bounding box. Finally, for all defined quadrants (each contain-

ing less than 5,000 photos) separate API queries can be executed. In this way, selected metadata such as user ID, image accuracy, user contact count, number of photos per user, and tag count per photo were downloaded (from the public photo pool) for photographs textually tagged as “Brandenburg Gate” and “Berlin” as well as “Reichstag” and “Berlin”.

The retrieved metadata for images for the POI are further filtered based on the scale at which the images were geotagged. This scale is called *accuracy* in Flickr and is derived from the zoom level of the map. The accuracy varies between 1 and 16, with 1 being at the world level and 16 being at the street level and representing the highest accuracy in Flickr. We extracted the metadata for Flickr images which have been geotagged at street level.

The retrieved geotags of the images are considered as observer points from where the photographs were taken. For the reverse viewshed calculation we use a Digital Surface Model (DSM) that represents the earth’s surface, including the elevation of buildings as well as the heights of the surrounding vegetation in our area of interest. These surface heights are derived from IRS-P5 Cartosat-1 in-flight stereo data with a 5 m post spacing and a relative vertical accuracy of 2.5 m with a linear error of 90% (LE90).

With the help of the surface creation tool in the Spatial Analyst toolbox in Esri’s ArcGIS 10.1 suite, (<http://www.esri.com/software/arcgis/arcgis10>), we computed multiple viewsheds from each observer point pertaining to each geotag of the Flickr images. For this study, we took a sample of 200 images, 100 for each POI. For each of those images, a viewshed was calculated. Afterwards, we analyzed for each image whether the calculated area of visibility includes the position of the POI (Brandenburg Gate or Reichstag). If that is the case, the image is considered to be *correctly geotagged* (Figures 3 and 4, green polygons). If the image content represents the POI, the image is also considered as *correctly labelled*. If the area of visibility does not include the position of the POI, the image is considered *incorrectly geotagged*

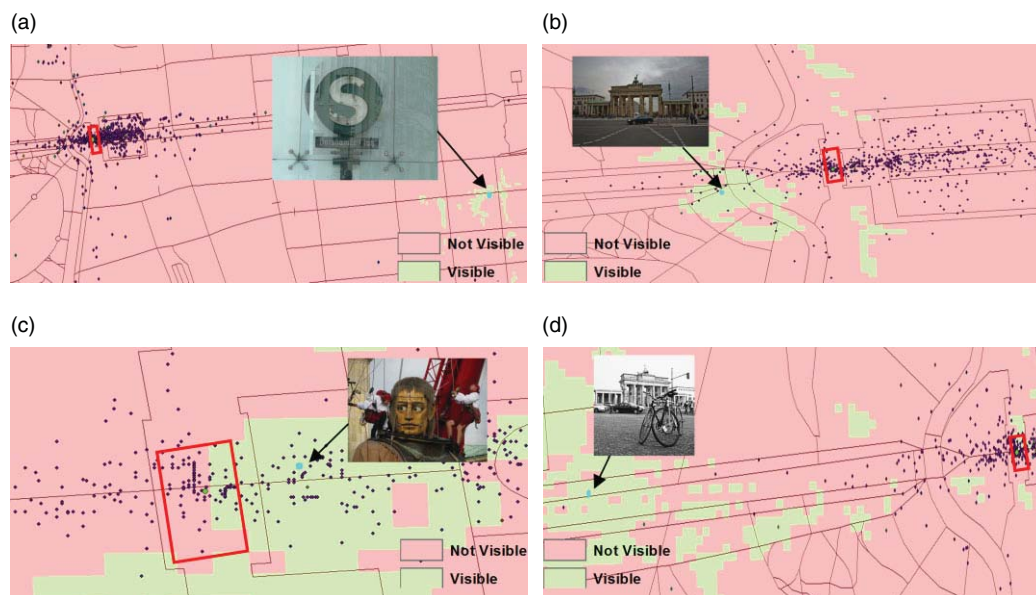


Figure 3 Left to right: a, b, c, d. The areas of visibility (green) from four different observer points to the Brandenburg Gate in Berlin (highlighted with red rectangle). The arrow points to the observer point and the image taken from there

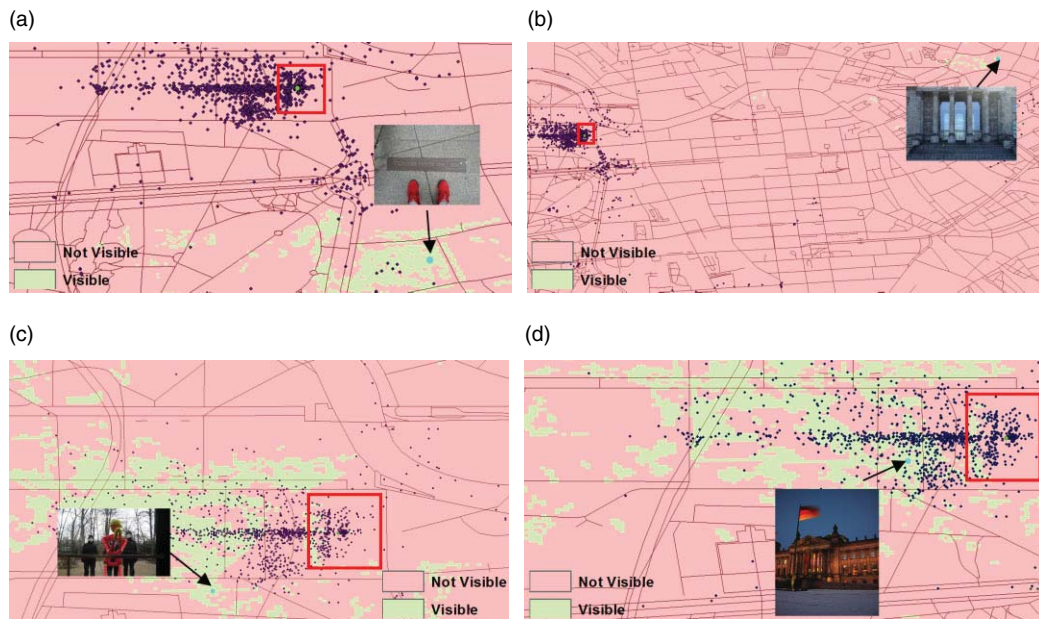


Figure 4 Left to right: a, b, c, d. The areas of visibility (green) from four different observer points to the Reichstag in Berlin (highlighted with red rectangle). The arrow points to the observer point and the image taken from there

(Figures 3 and 4, pink polygons), as the observer could not have seen the POI. If the image content does not represent the POI, it is considered as *incorrectly labelled*. Those considerations result in four different categories an image can belong to: (a) images incorrectly geotagged and incorrectly labelled; (b) images incorrectly geotagged, but correctly labelled; (c) images correctly geotagged, but incorrectly labelled; and (d) images correctly geotagged and correctly labelled. These four categories within the Brandenburg Gate and the Reichstag use cases are depicted in Figures 3a-d and 4a-d. It should also be noted here, that photographs that were taken from elevated locations such as a higher floor of a building are disregarded in our analysis, as the height of the position with which the photograph was taken is not included in the viewshed computation.

Photographs that are geotagged out of the visibility range (POI falls in pink coloured areas) are either considered to misrepresent the location from where the photograph was taken, or represent something other than the POI although tagged as the latter. Photographs belonging to either of these two categories are considered to represent an incorrect location for the POI.

4 Analysis of Visual VGI Metadata for User Credibility Assessment

Next, we explore how we can build on the described approach for assessing the location correctness of visual VGI, towards inferring the credibility of VGI users. We propose here to analyze the dependency relationship between metadata attributes (e.g. user contacts count) and the location correctness of the geotags; i.e., we determine the location correctness of Flickr

Table 1 The categories of images within the sample dataset falling into correct/incorrect geotagging and labelling

Category	Correct Geotag	Correct Label
<i>a</i>	No	No
<i>b</i>	No	Yes
<i>c</i>	Yes	No
<i>d</i>	Yes	Yes

geotags through the reverse viewshed analysis, consider it as an example of a reference quality measurement for Flickr photographs, and relate it to user and photo metadata attributes. Related research such as Zwol (2007), Castillo et al. (2011), or Gupta et al. (2012) utilized various VGI user metadata to derive conclusions and to characterize the user. Zwol (2007) takes the number of contacts of a user as the predictor for the expected popularity of a photo within the Flickr data source. Therefore, it can be assumed that the user contact number characterises to a certain degree the popularity of the user. Further, Castillo et al. (2011) and Gupta et al. (2012) showed for Twitter data how user-based features, such as the user friend count and contribution frequency, associate with *information* credibility. This shows that user profile features can be used as a rich source of information to derive characteristics about the user, including content credibility.

Based on these works, and in combination with the reverse viewshed as a reference quality measure, we can explore which user metadata shows a pattern within users who correctly and incorrectly geotag a photograph. For each of the two selected points of interest, the Brandenburg Gate and Reichstag in Berlin, we analyzed 100 geotagged Flickr images, each for its image content together with its photo and user metadata. Our analysis is summarized in Table 1 and 2. The photos are classified as *a* (wrong geotag and wrong label), *b* (wrong geotag but correct label), *c* (correct geotag but incorrect label) and *d* (correct geotag and correct label) (Table 1).

Table 2 presents the variation of each metadata attribute within the four image categories *a*, *b*, *c*, and *d* for Brandenburg Gate and Reichstag. To complement Table 2, Figures 5 to 12 present the descriptive statistics of the selected metadata elements for the four identified categories. We can observe interesting patterns within the gathered data. Producers of photos within category *b* and *d* for both POI have on average the lowest number of contacts (on average 121 contacts for “Brandenburg Gate” images and 125 contacts for “Reichstag” images), as compared with producers of photos with incorrect labels (categories *a* and *c*) who have on average 236 contacts within “Brandenburg Gate” images and 130 contacts within “Reichstag” images. This may explain the motivation and thus different priorities of users when contributing to VGI as also described by Coleman et al. (2009). Users who have correctly labelled their images tend to have on average fewer contacts than users falling into the remaining categories. Hence, popularity in Flickr may not be a priority for this group of users, while quality is a priority.

Furthermore, we looked into the average number of photos contributed by users to Flickr within each category. This also revealed a pattern of correct and incorrect image labelling. Producers of photos of categories *a* and *c*, with incorrect labels, have contributed significantly more photos over the years of their participation on Flickr. The average photo count of photo producers for POI Brandenburg Gate in category *a* is 19,087 and for category *c* is 18,354,

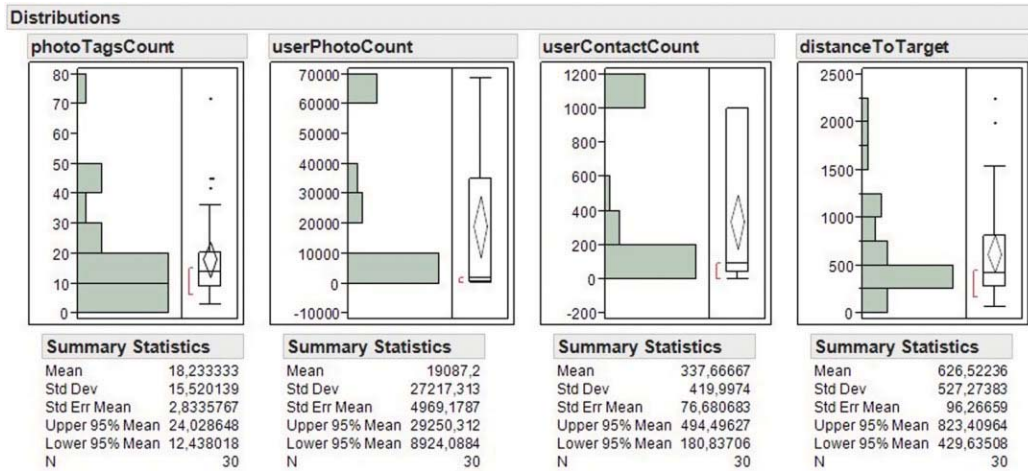


Figure 5 Distribution of data for category 'a' within the Brandenburg Gate use case

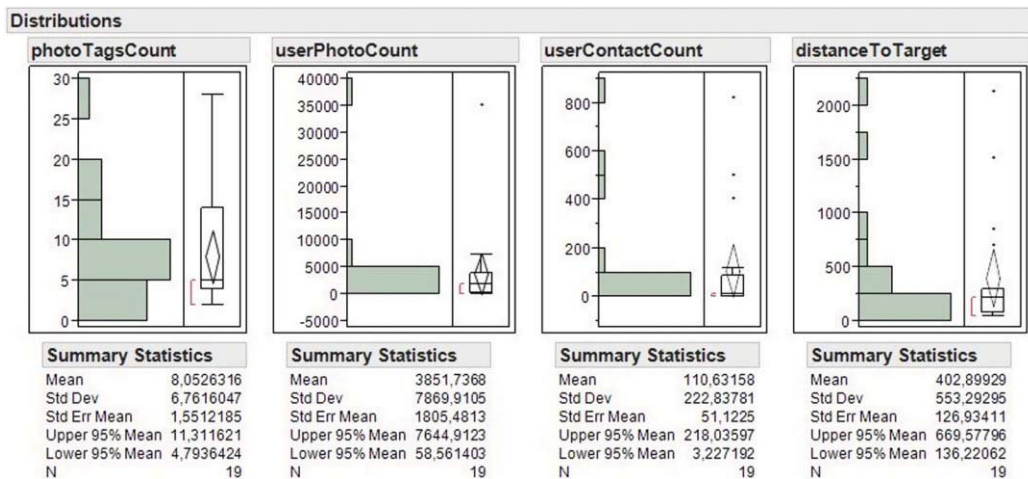


Figure 6 Distribution of data for category 'b' within the Brandenburg Gate use case

while for category *d* it is 5,422 and for category *b* it is 3,852. The average photo count of photo producers for POI Reichstag in category *a* is 8,136, category *c* is 9,555 while for category *b* and *d* it is 7,928 and 2,618, respectively.

Looking into the photo metadata, the average number of tags per photo further reveals a pattern in the above image categories. Photos for the Brandenburg Gate within categories *a* (18 tags on average) and *c* (13 tags on average) have on average the highest number of tags. These photos are incorrectly labelled. Whereas photos in category *b* (8 tags on average) and *d* (11 tags on average) have the lowest number of tags on average and are also correctly labelled. Likewise, photos for the Reichstag within categories *a* (35 tags on average) and *c* (22 tags on average) have on average the highest number of tags per photo, and photos in category *b* (12 tags on average) and *d* (10 tags on average) have the lowest number of tags on average and are also correctly labelled.

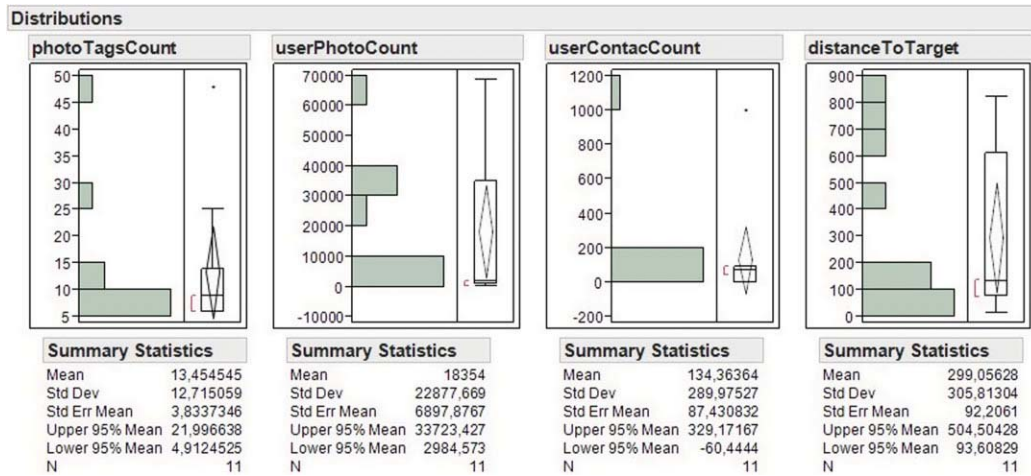


Figure 7 Distribution of data for category 'c' within the Brandenburg Gate use case

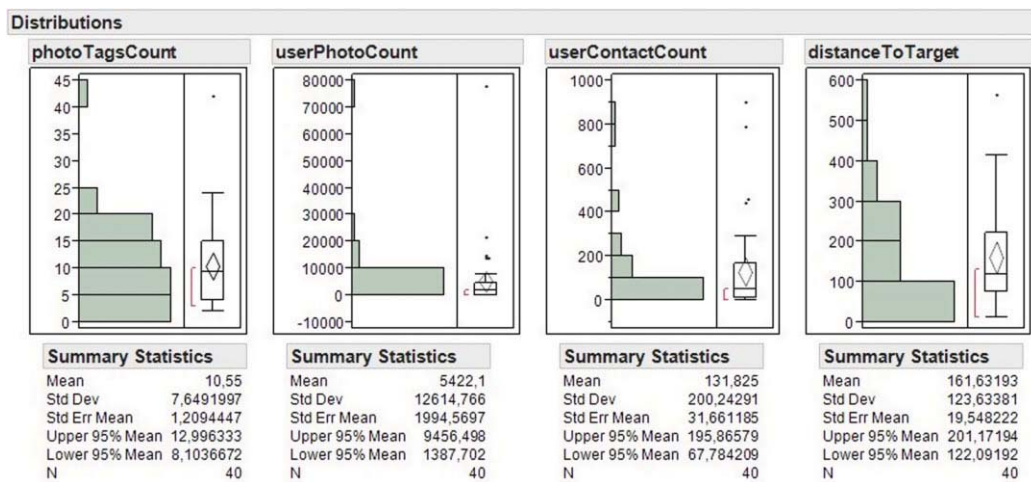


Figure 8 Distribution of data for category 'd' within the Brandenburg Gate use case

Further, we have computed the distance to the target by taking the orthodrome between the geotag and the actual geographical coordinates of a POI. This reveals that the average distance to the target decreases for images from *a* to *d* within the use cases for the Brandenburg Gate as well as the Reichstag. Images in category *a* have the longest average distance to the target and in category *d* have the shortest average distance to the target (Table 2). The closer to the POI a person is, the more focused the object would be in the image, thus allowing the person to geotag/label more precisely. The further away from the POI the person is, he or she might become less precise when geotagging and labelling the image.

The above observations can be considered as triggers to look further into these findings. They will enable us to infer the user credibility within similar VGI sources, and in general to understand qualitative aspects in user provided data much better. In addition to the location

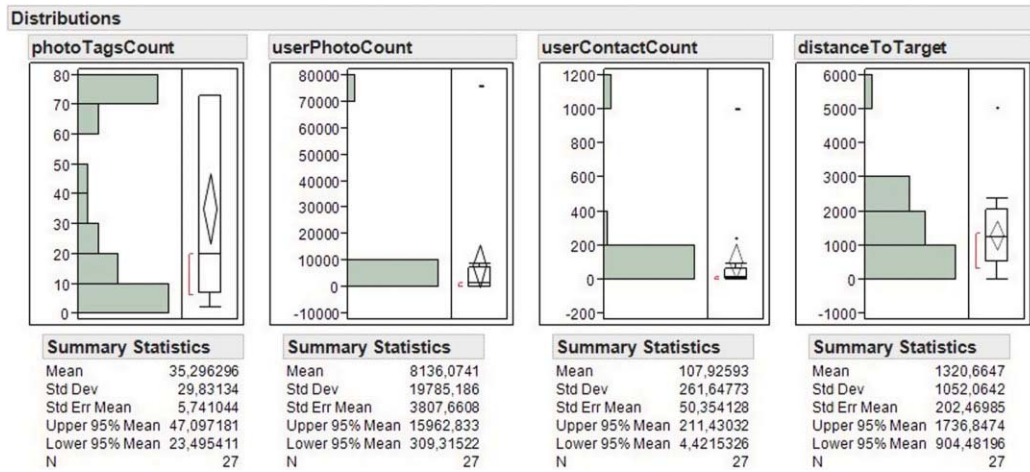


Figure 9 Distribution of data for category 'a' within the Reichstag use case

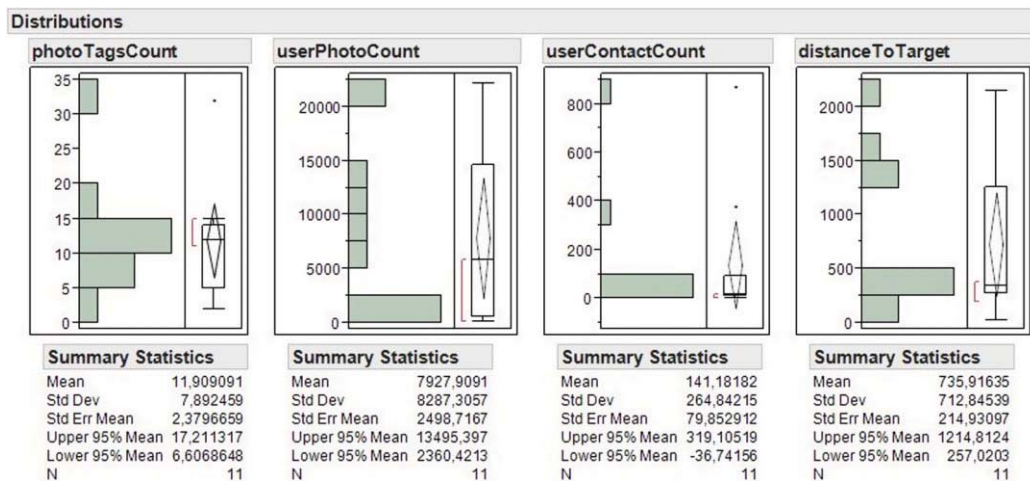


Figure 10 Distribution of data for category 'b' within the Reichstag use case

correctness, other features such as the label precision or image content can be used to evaluate user credibility. Methods for utilizing these features in combination to assess user credibility are discussed in the following sections.

5 Discussion

A reverse viewshed is carried out to assess the location correctness of geotagged Flickr images that confirm to a particular POI through the geotag and the image label. Images placed within a visibility region that do not include the position of the POI are determined to be either incorrectly geotagged, incorrectly labelled, or both. We have to consider possible reasons for these outliers. An obvious reason are mistakes made by the user when geotagging a photo. Such

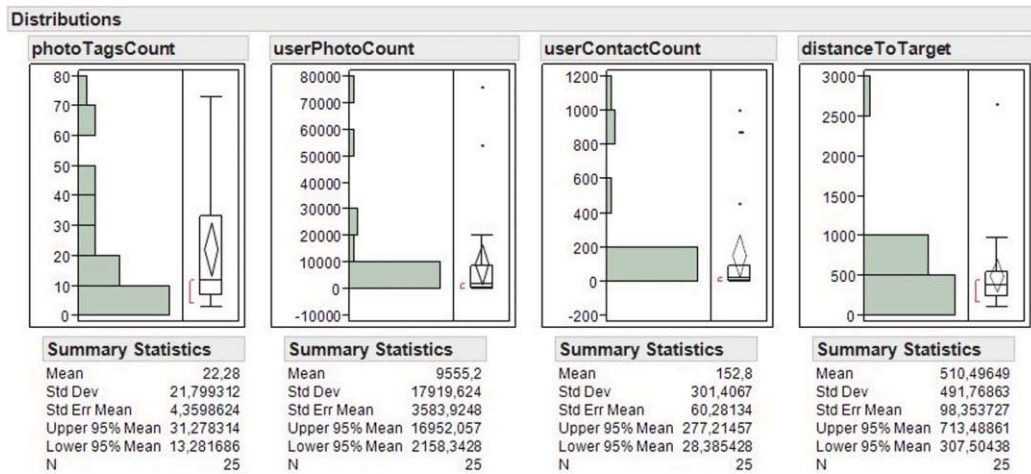


Figure 11 Distribution of data for category 'c' within the Reichstag use case

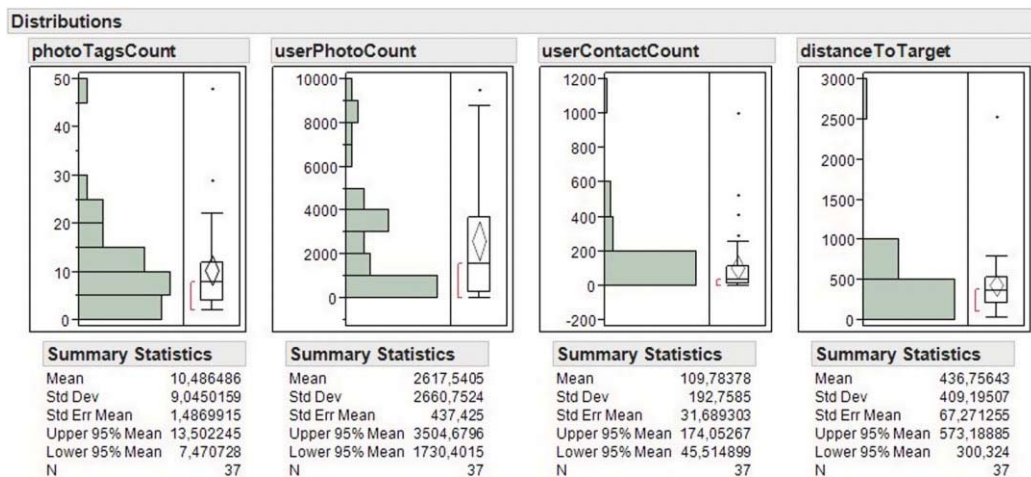


Figure 12 Distribution of data for category 'd' within the Reichstag use case

mistakes can for example result from either manually adding, as a geotag to the photo, wrongly measured coordinates or coordinates measured by a malfunctioning GPS device. Another reason might be that, while the geotag is correct, the user lacks sufficient knowledge about what is shown on the photograph and provides incorrect place-description tags. In addition, we have seen cases within the data sets, where it seems that users have made touristic round trips and collectively tagged their taken photos with all places visited during that trip. For example, a tourist visiting several places in Germany defines the same tags (including "Brandenburg Gate") for all taken photos during his/her trip and bulk uploads them as a photo set to Flickr.

The cases above can be clearly considered as wrongly tagged photos; and lowering the credibility of the producers of such photos would be valid. Other outliers cannot be as easily considered as being wrongly tagged. In particular, when extracting data for a particular place

of interest based on their textual tagging, we have to encounter outliers that are duplicates and referred to by the same name. One such example is the Eiffel Tower replica in Las Vegas (a replica of the original in Paris), which also attracts many visitors. Another example are photos that show miniatures of important sights. They are validly tagged by a user with the name of that sight while being located far away from the original place of interest. An example is the photo of a miniature Eiffel Tower on someone's desk. A difficult case is a collection of photos of a certain place where a user makes comparisons with other sights by adding the compared place of interest as a tag. An example could be a photo of the Shibuya crossing in Tokyo where the producer wants to point out that it looks similar to Times Square in New York and provides tags accordingly. Hence, a complement to our approach would be to utilize image recognition techniques that can programmatically identify the image content and compare it with the POI to find (dis)similarities, and then associate it with the reverse viewshed. This would filter out images that are irrelevant to our query (e.g. those that are textually/geographically tagged as the Brandenburg Gate but represent a bus stop in the nearby region), and show us images that represent the target within the reverse viewshed. Text analysis algorithms can also help us in filtering out relevant and irrelevant labelled photographs.

Thus far, we have considered only one aspect with which the reliability of a photograph can be assessed: the location correctness. In addition to this there are further aspects, as described above, that contribute to the reliability of an image, such as the label completeness, content relevance, user profile completeness, etc. A weighted score for each of these aspects could give us a complete reliability score for each user, with which the user credibility can be evaluated.

Regarding data accuracy, when computing the (reverse) viewshed analysis, one has to encounter issues of output quality variability that were emphasized by Fisher (1991). Those quality issues are due to data errors, data resolution, and errors in the viewshed analysis algorithm. Thus, within this article we limit our approach to calculating a reverse viewshed upon which the location correctness of geotagged Flickr images are assessed. We propose to use additional user/photo metadata in combination with the location correctness to infer the credibility of users as an extension to future work.

6 Conclusions and Future Work

This article contributes to the research and discussion on quality control of VGI. We have investigated through experimental analysis how a reverse viewshed analysis can be utilized to assess the location correctness of visually generated VGI. In doing so, we have first programmatically downloaded metadata of photographs for a certain POI by querying the open Flickr API for all geotagged photos that are textually tagged (labelled) with the place description (e.g. with the tags "Brandenburg Gate" and "Berlin"). As a next step, we have computed the area of visibility from each observer point (geotag) based on surface elevation data, to the given POI, the Brandenburg Gate and the Reichstag in Berlin. With the help of this reverse viewshed analysis we were able to determine if the position of the POI lies within the visibility of a given observer point. If it lies outside the visibility region, the photograph captured by the observer is considered as incorrectly geotagged. We duly note that all images that do correspond to the POI through the geo/text tag do not necessarily visually represent the POI. This is also exhibited through analyzing a sample dataset. We propose in future work to conduct image recognition techniques to filter out images that are irrelevant to the POI.

Within the sample dataset for Brandenburg Gate and Reichstag we have categorized the photographs into four groups based on the geotag and label correctness. In those categories we made observations about user and photo metadata. In particular, we have found that users producing photos for category *a* and *c* (both wrongly labelled) have on average higher numbers of photos (for both use cases) and higher numbers of tags. Further, the producers of photos in category *b* and *d* (correctly labelled) together have on average a lower number of contacts than the other photo categories. As we insinuate that these are valuable indications for assessing the credibility of users based on the reliability of their contributions, these further imply investigating the tagging behavior of users beyond their motivational aspects.

For the future, we will work towards a mechanism for automatically inferring user credibility through analyzing the dependency between certain user metadata and the reference quality measure, the location correctness determined with the reverse viewshed. Thereby, the influence of viewshed sensibility will be studied and optimized, e.g. by investigating vectorized city models based on CityGML. Further, we will look into the possibility of extracting credibility-related measures from analyzing free-text comments that users provide for photos. An example is sentiment analysis, which computes polarity scores regarding the expressed opinions. Another direction will be to look into the temporal trends of photo capturing and uploading behavior. Looking into these additional aspects and giving them a weighted score to find the complete reliability of geotagged images will allow us to evaluate the user's credibility within these visually generated VGI sources. Furthermore, to automate the process of user credibility assessment we can envisage training statistical prediction algorithms for classifying the users according to the abovementioned weighted reliability parameters.

References

- Ames M and Naaman M 2007 Why we tag: Motivations for annotation in mobile and online media. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*, San Jose, California: 971–80
- Becker C and Bizer C 2011 Flickr TM Wrapp: Precise Photo Association. WWW document, <http://wifo5-03.informatik.uni-mannheim.de/flickrwrapp/>
- Bishr M and Kuhn W 2007 Geospatial information bottom up: A matter of trust and semantics. In Fabrikant S I and Wachowicz M (eds) *The European Information Society*. Berlin, Springer Lectures Notes in Geoinformation and Cartography: 365–87
- Castillo C, Mendoza M, and Poblete B 2011 Information credibility on Twitter. In *Proceedings of the Twentieth ACM International Conference on the World Wide Web*, Hyderabad, India: 675–84
- Coleman D J, Geogiadou Y, Labonte J 2009 Volunteered geographic information: The nature and motivation of producers. *International Journal of Spatial Data Infrastructures Research* 4: 332–58
- Ciepluch B, Jacob R, Mooney P, and Winstanley A 2010 Comparison of the accuracy of OpenStreetMap for Ireland with Google Maps and Bing Maps. In *Proceedings of the Ninth International Symposium on Spatial Accuracy Assessment in Natural Resources and Environmental Sciences*, Leicester, United Kingdom
- Crandall D J, Backstrom L, Huttenlocher D, and Kleinberg J 2009 Mapping the world's photos. In *Proceedings of the Eighteenth International Conference on the World Wide Web (WWW '09)*, Madrid, Spain: 761–70
- De Florian L, Marzano L, and Puppo P E 1994 Line-of-sight communication on terrain models. *International Journal of Geographic Information Systems* 8: 329–42
- Flanagin A J and Metzger M J 2008 The credibility of volunteered geographic information. *GeoJournal* 72: 137–48
- Friedland G, Choi J, Lei H, and Janin A 2011 Multimodal location estimation on Flickr videos. In *Proceedings of the Nineteenth ACM International Conference on Multimedia*, Scotsdale, Arizona
- Frew J 2007 Provenance and Volunteered Geographic Information. WWW document, http://www.ncgia.ucsb.edu/projects/vgi/docs/position/Frew_paper.pdf
- Fisher P F 1991 First experiments in viewshed uncertainty: The accuracy of the viewshed area. *Photogrammetric Engineering and Remote Sensing* 57: 1321–27

- Fisher P F 1993 Algorithm and implementation uncertainty in viewshed analysis. *International Journal of Geographic Information Systems* 7: 331–47
- Fisher P F 1996 Extending the applicability of viewsheds in landscape planning. *Photogrammetric Engineering and Remote Sensing* 62: 1297–1302
- Girardin F, Blat J, Calabrese F, Fiore F D, Ratti C 2008 Digital Footprinting: Uncovering Tourists with User-generated Content. *IEEE Pervasive Computing* 7: 36–43
- Girres J F and Touya G 2010 Quality assessment of the French OpenStreetMap dataset. *Transactions in GIS* 14: 435–59
- Golder S and Huberman B 2006 Usage patterns of collaborative tagging systems. *Journal of Information Science* 32: 198–208
- Goodchild M F 2007 Citizens as sensors: The world of volunteered geography. *GeoJournal* 69: 211–21
- Goodchild M F 2009 NeoGeography and the nature of geographic expertise. *Journal of Location Based Services* 3: 82–96
- Goodchild M F and Li L 2012 Assuring the quality of volunteered geographic information. *Spatial Statistics* 1: 110–20
- Gupta M, Zhao P, and Han J 2012 Evaluating event credibility on Twitter. In *Proceedings of the SIAM International Conference on Data Mining (SDM 2012)*, Anaheim, California
- Haklay M 2010 How good is volunteered geographical information? A comparative study of OpenStreetMap and Ordnance Survey datasets. *Environment and Planning B* 37: 682–703
- Haklay M, Basiouka S, Antoniou V, and Ather A 2010 How many volunteers does it take to map an area well? The validity of Linus' law to volunteered geographic information. *Cartographic Journal* 47: 315–22
- Hollenstein L and Purves R 2011 Exploring place through user generated content: Using Flickr tags to describe city cores. *Journal of Spatial Information Science* 1: 21–48
- Hovland C I, Janis I L, and Kelley J J 1953 *Communication and Persuasion*. New Haven, CT, Yale University Press
- Jankowski P, Andrienko N, Andrienko G, and Kisilevich S 2010 Discovering landmark preferences and movement patterns from photo postings. *Transactions in GIS* 14: 833–52
- Kidner D, Sparkes A, and Dorey M 1999 GIS and wind farm planning. In Stilwill J, Geertman S, and Openshaw S (eds) *Geographical Information and Planning*. London, Springer: 203–23
- Kim Y, Rana S, and Wise S 2004 Exploring multiple viewshed analysis using terrain features and optimisation techniques. *Computers and Geosciences* 30: 1019–32
- Lake I R, Lovett A A, Bateman I J, and Langford I H 1998 Modelling environmental influences on property prices in an urban environment. *Computers, Environment and Urban Systems* 22: 121–36
- Montello D R, Goodchild M F, Gottsegen J, and Fohl P 2003 Where's downtown? Behavioral methods for determining referents of vague spatial queries. *Spatial Cognition and Computation* 3: 185–204
- Moxley E, Kleban J, and Manjunath B S 2008 Spirittagger: A geo-aware tag suggestion tool mined from Flickr. In *Proceedings of the First ACM International Conference on Multimedia Information Retrieval*, Vancouver, British Columbia: 24–30
- O'Sullivan D and Unwin D J 2003 *Geographic Information Analysis*. New York, John Wiley and Sons
- Ralling P, Kidner D, and Ware A 1999 Distributed viewshed analysis for planning applications. In Getting B (ed) *Innovations in GIS* 6. London, Taylor and Francis: 185–99
- Rinner C, Kessler C, and Andrulis S 2008 The use of Web 2.0 concepts to support deliberation in spatial decision-making. *Computers, Environment and Urban Systems* 32: 386–95
- Samet H 1984 The quadtree and related hierarchical data structures. *ACM Computer Surveys* 16: 187–260
- Sherren K, Fischer J, Pink J, Stott J, Stein J, and Yoon H 2011 Australian graziers value sparse trees in their pastures: A viewshed analysis of photo-elicitation. *Society and Natural Resources* 24: 412–22
- Sigurbjoernsson B and Van Zwol R 2008 Flickr tag recommendation based on collective knowledge. In *Proceedings of the Seventeenth International Conference on the World Wide Web (WWW '08)*, Beijing, China: 327–36
- Surowiecki J 2005 *The Wisdom of Crowds*. New York, Anchor
- Valli C and Hannay P 2010 Geotagging where cyberspace comes to your place. In *Proceedings of the International Conference on Security and Management (SAM '12)*, Las Vegas, Nevada: 627–32
- Wang J, Robinson G J, and White K 1996 A fast solution to local viewshed computation using grid-based digital elevation models. *Photogrammetric Engineering and Remote Sensing* 62: 1157–64
- Zwol R V 2007 Flickr: Who is looking? In *Proceedings of the IEEE/WIC/ACM International Conference on Web Intelligence*, Silicon Valley, California