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

Hansi Senaratne[1], Arne Bröring[2, 3], Tobias Schreck[1]

[1] University of Konstanz, Germany

[2] Institute for Geoinformatics, University of Münster, Germany

[3] 52°North Initiative for Geospatial Open Source Software GmbH, Münster, Germany

ABSTRACT

With the increased availability of user generated data, assessing quality and credibility of such data becomes important. In this paper, we propose to assess the location correctness of visually generated 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 utilise Flickr photographs. Therefore, we first collect all Flickr photographs that confirm 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 from the observer points. If the point of interest lies outside the visibility from a given observer point, the respective geotagged image is considered to be incorrectly geotagged. This way, we analyse 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. E.g., attributes such as profile completeness together with the location correctness can serve as a weighted score to assess credibility.

Address for correspondence:

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

1. INTRODUCTION

In today's information-driven society, Volunteered Geographic Information (VGI) has shown an immense increase over the past years. With this massive increase of data production by volunteers, the need for cautiousness towards data credibility becomes even 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 conceptualisation of location is due not only to the fact that geographic entities are continuous in nature, but also due to the quality and limitations of spatial knowledge (Hollenstein & Purves, 2011).

Hovland et al. (1953) expressed *credibility* as the believability of a source or message, which comprises primarily of two dimensions, the *trustworthiness* and *expertise*. Flanagin & Metzger (2008) further asserted that, while trust and expertise have different meaning from credibility as well as differing meaning between each other, one conceives 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 that contribute to this perception of trustworthiness and believability, other than data accuracy itself. Metadata about the origin of VGI can provide a foundation for judgment on the quality and trust (Frew, 2007).

In case of Flickr, as an example of a platform for visually generated VGI, volunteers can upload photographs to share them with others. A Flickr user can maintain a profile to which uploaded photos are linked and to 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 form of geographic coordinates. This *geotag* can be either 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., neighbourhood, city, country).

Additionally to the geotag that consists of geographic coordinates, Flickr users often specify the place of interest to which the picture relates, as textual tags. 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 the ancient city in Cambodia, this visual analysis shows that there are also many pictures being located far away from it[1]. Also Becker & Bizer (2011) demonstrated through their work on the *Flickr Wrappr*, how pictures on Flickr are incorrectly geotagged.

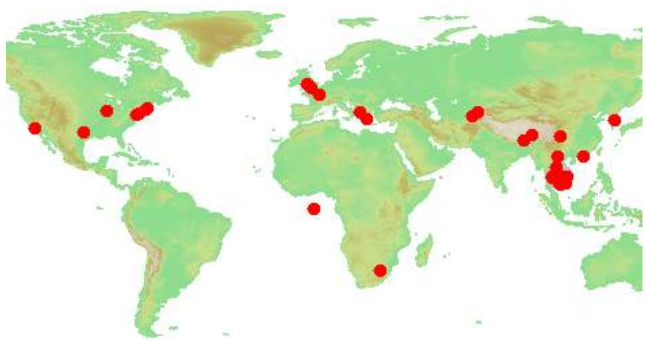

Figure 1. Geotags of Flickr photos that were textually tagged as "Angkor" and "Cambodia".

In this paper, 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 generalises 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 pints of interest (POIs), 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 POIs is calculated. A reverse viewshed analysis holds the same

---

[1] For example, this photo displaying a part of Angkor is geotagged at a location in California:

http://www.flickr.com/photos/rbleib/5030263322/in/set-72157624911484519/

principles as the viewshed analysis, however, it is utilised to determine the visibility to a given point of interest (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 the image is considered to either misrepresent the location from where the photograph was taken, or the photographed content represents something else other than the POI but tagged as the latter. Photographs belonging to either of these two groups are considered to be tagged with incorrect location.

Using this approach, we investigate which metadata of photographs (e.g., tag count of photographs, comments count of photographs, etc.) as well as metadata about users (e.g., the number of photos, the number of contacts, or the used camera) can be utilised to eventually infer the credibility of photographers regarding a correct geotagging. We achieve this through analysing 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 and also to transfer this methodology to other visual VGI sources (e.g. Panoramio).

The remainder of this paper is organised as follows. In Section 2 we review selected related work to our research. In Section 3 we discuss our approach in detail. Section 4 analyses 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 & 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). Surowieki (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[2], Wikimapia[3], Google MyMaps[4], Flickr[5]). Goodchild (2007) coined this phenomenon as Volunteered Geographic Information (VGI). As of January 2012, Flickr has reported to host over 6 billion images[6] (around 3% of the Flickr images were geotagged[7] in 2009), and OpenStreetMap statistics[8] state that over a million registered users have contributed more than 3 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 & 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 done on assessing the credibility of user generated geo information on different platforms. Thereby, the central question has been, whether we can trust the data volunteers to produce data of *usable* quality, which suffice for convenient usage of the data to derive accurate conclusions. A few VGI platforms have taken measures to moderate the credibility of user generated data

---

[2] www.openstreetmap.org

[3] www.openstreetmap.org

[4] https://www.google.com/maps/mm

[5] www.flickr.com

[6] http://www.searchenginejournal.com/the-growth-of-social-media-an-infographic/32788/

[7] http://code.flickr.net/2009/02/04/100000000-geotagged-photos-plus/

[8] http://www.openstreetmap.org/stats/data_stats.html

to ensure data reliability. For example, the Audubon Society's Christmas Bird Count[9] is open for bird watchers to observe migration patterns, bird population etc., and to contribute these observations on the open platform. In order to take part, the users should possess domain knowledge to a certain degree and they are also given proper training before they can take part. Project GLOBE[10] is another example that 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 guide the students throughout the process. WikiScanner[11] is an example for assessing *user* credibility in UGC. WikiScanner cross references the edits on Wikipedia[12] with the data on the editor of the associated block of IP addresses of various organisations[13]. This author identity is what provides credibility judgment.

Extending the work by Haklay (2010), Girres & Touya (2010) assessed the quality of OpenStreetMap data for France by comparing them to 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 quality of OpenStreetMap data, Haklay et al. (2010) found that positional accuracy of features improves as the number of editors increases up to 13. Goodchild & Li (2012) proposed a three tier approach to assuring VGI quality: i) Crowd-sourcing (number of contribution and accuracy), ii) volunteers who are given roles in the *hierarchy* to moderate the

---

[9] http://birds.audubon.org/christmas-bird-count

[10] http://training.globe.gov/

[11] http://wikiscanner.virgil.gr/

[12] www.wikipedia.org

[13] Although, Wikiscanner does not distinguish between edits made by authorised users from IP addresses originating from organisations and edits made by unauthorised intruders and users of public access computers.

data accuracy, and iii) the Geographic approach where geographic features on a map are inferred from knowledge on 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 & 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 categorise 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. I.e., there is no clear role of the volunteer belonging to any one of these distinguished roles. However in contributing to VGI, they are all experts in their own local communities. On the contrary, volunteers falling into academic geography are involved in professional geography, such as surveyor or cartographer. Coleman et al. (2009) classified data volunteers as overlapping between *Neophytes*, *Interested Amateur*, *Expert Amateur*, *Expert Professional*, and *Expert Authority*. He analysed these groups based on 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[14] or Garmin[15].

[14] http://www.tomtom.com/

Volunteers falling into *Social Networks* contribute to sources such as OpenStreetMap, Flickr etc. Volunteers falling into Civic/Governmental contribute data out of concern to their city/society, for example to PPGIS[16].

Zwol (2007) presents a characterisation of user behaviour 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 which is sustained by the network of contacts within Flickr, is important for the popularity of their photos. In other Flickr analyses, Friedland et al. (2011) as well as Moxley et al. (2008) utilise 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 Behaviour in Flickr

Flickr photos have been explored in a multitude of geographical analyses. For instance, Jankowski et al. (2010) and Crandall et al. (2009) explored spatial and temporal patterns in user movement and their interests in landmark and events captured through Flickr. These Flickr photos are organised or searched with the help of their accompanying tags that come in various forms. Ames & Naaman (2007) have comprehensively discussed the concept of *tagging* and have identified two main incentives that motivate users to tag: i) *sociality*, describing who is intended to use the tag, ii) *function*, describing the intended usage of the tag, which could be either for organisational 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 form of specific keywords to describe the content (Golder & Huberman, 2006), or in form of geographic coordinates (Geotagging) to identify the location linked to the image content (Valli & Hannai, 2010). Moxley et al. (2008) developed a tool that suggest tags for a given image, based on the geographic context and visual

---

[15] http://www.garmin.com/us/

[16] http://www.ppgis.net/

relevance. Crandall et al. (2009) analyse 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 to estimate 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) analysed tags of Flickr photos to explore how people perceive their environment and the underlying semantics on how they describe the urban space, Sigurbjoernsson & 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 on what *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 geo coordinates 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 evaluate the location correctness of visually generated VGI (Section 3.1). Section 3.2 and 3.3 describe in detail our methodology, 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 POIs, i.e., Brandenburg Gate and Reichstag in Berlin, are automatically downloaded by a crawler. The label of the POIs 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 point of interest. A reverse viewshed successively determines from which observer points the point of interest is visible. This allows cross validating if a picture was taken within the vicinity to the point of interest. 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.

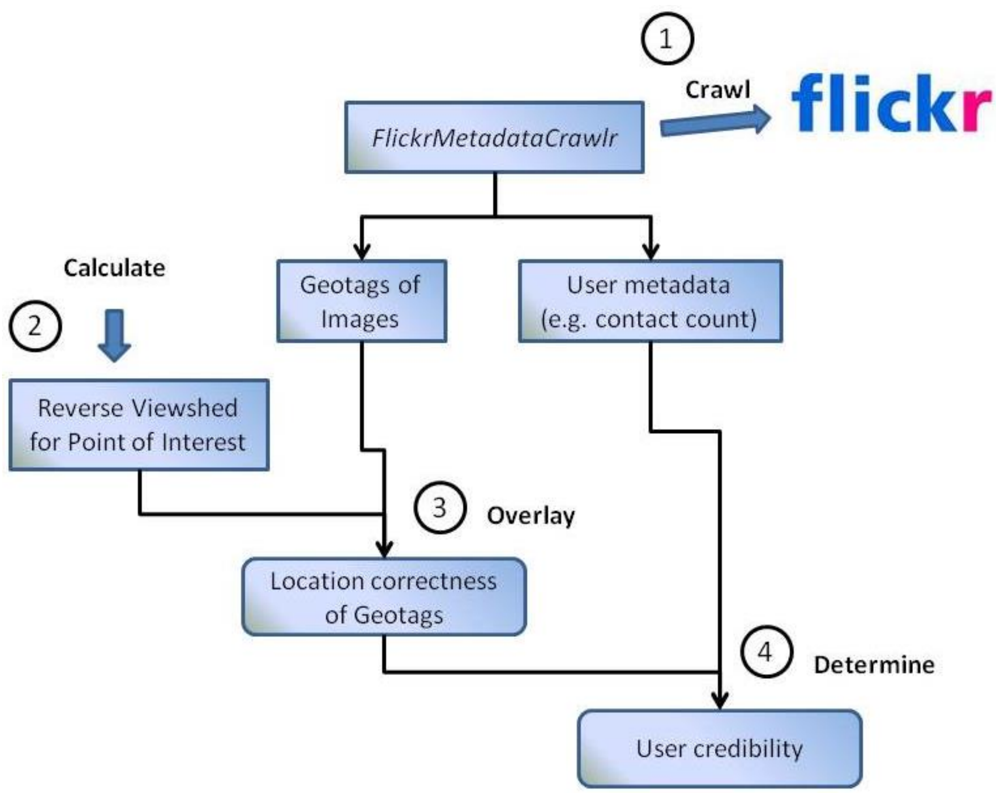

Figure 2. The workflow of the proposed methodology to assess the location correctness of visually generated VGI. We utilise Flickr as the data source.

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 POIs.

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 & 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. 2010). 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 analysed 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). Amongst 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 utilised 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 utilised the same technique to generate a viewshed but a different procedure. I.e., instead of taking one viewshed from the target point, we create multiple viewsheds from the observer points to validate if 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[17]. This tool is able to programmatically download metadata of Flickr photos and users. The FlickrMetaCrawlr therefore relies on the open Flickr API[18] 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 utilises 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 of 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 containing less than 5,000 photos) separate API queries can be executed. 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 POIs are further filtered based on the scale at which the images were geotagged. This scale is called *accuracy* in Flickr which is derived from the zoom level of the map. The accuracy varies between 1 and 16, while 1 being at the world level and 16 being at the street level

---

[17] The source code of our *FlickrMetaCrawlr* can be accessed here: http://ifgi.uni-muenster.de/~arneb/FlickrMetaCrawlr.jar

[18] http://www.flickr.com/services/api

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 manmade 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 5m post spacing and a relative vertical accuracy of 2.5m with linear error of 90% (LE90).

With the help of the surface creation tool in the spatial analyst toolbox in ESRI's ArcGIS 10.1[19] suite, 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 analysed 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* (Figure 3; 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; 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 use case are depicted in the following Figures 3a to 3d. It should also be noted here, that photographs that were taken from elevated location such as a higher floor of a

[19] http://www.esri.com/software/arcgis/arcgis10

building are disregarded in our analysis, as the height of the position with which the photograph was taken in not included in the viewshed computation.

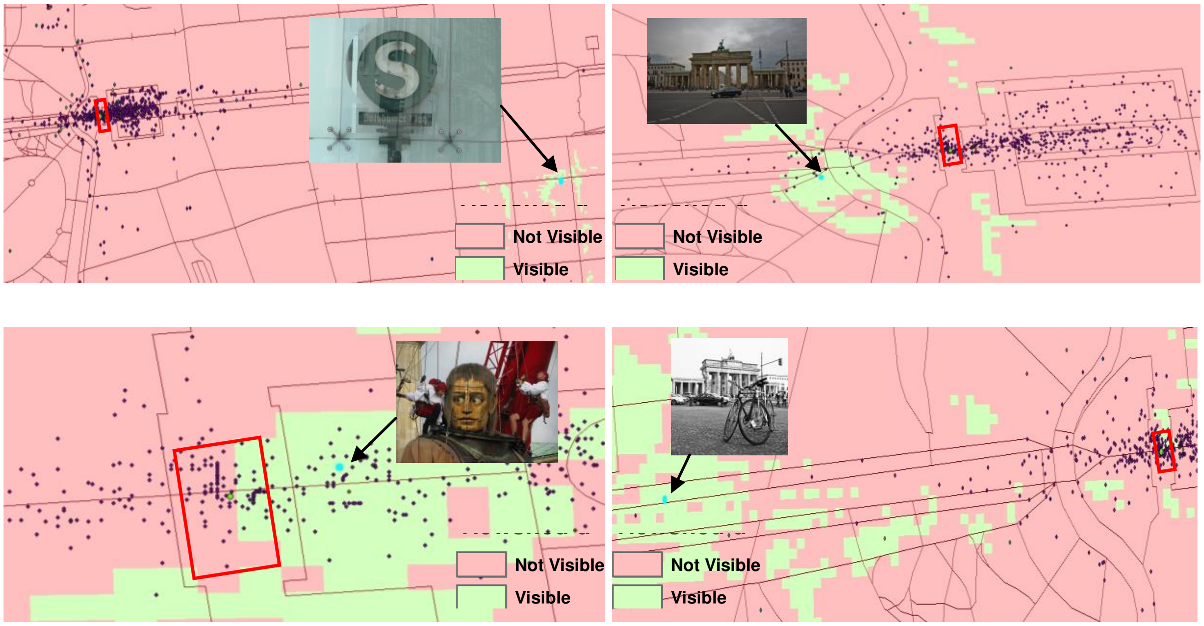

Figure 3. Left to right: a, b, c, d. The areas of visibility (green) from four different observer points to the Brandenburg Gate in Berlin (highlighted with red rectangle).The arrow points to the observer point and the image taken from there.

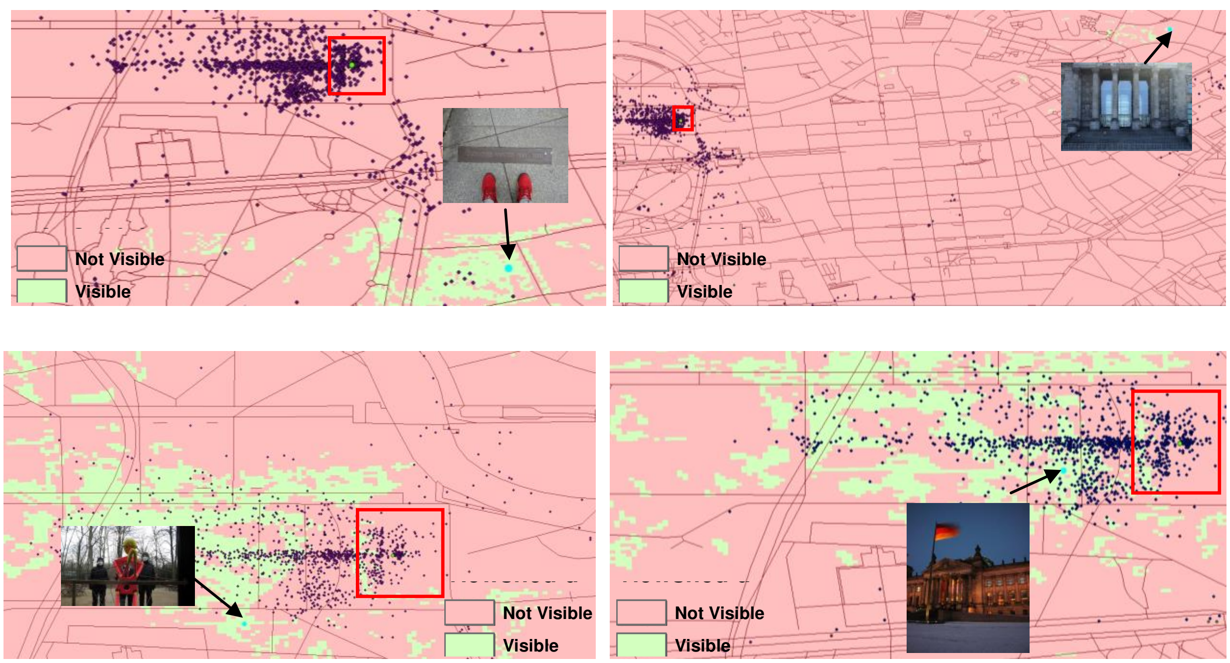

Figure 4. Left to right: a, b, c, d. The areas of visibility (green) from four different observer points to the Reichstag in Berlin (highlighted with red rectangle). The arrow points to the observer point and the image taken from there.

Photographs that are geotagged out of the visibility range (POI falls in pink coloured areas) are considered to either misrepresent the location from where the photograph was taken, or the photographed content represents something else other than the POI but tagged as the latter. Photographs belonging to either of these two are considered to be representing incorrect location of the point of interest.

4. Analysis of Visual VGI Metadata for User Credibility Assessment

Next, we explore how we can build up on the described approach for assessing the location correctness of visual VGI, towards inferring the credibility of VGI users. We propose here to analyse 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 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) utilised various VGI user metadata to derive conclusions and to characterise 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 analysed 100 geotagged Flickr images, each for its image content together with its photo and user metadata. Our analysis is summarised 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 1. The categories of images within the sample dataset falling into correct/incorrect geotagging and labelling.

Category	Correct Geotag	Correct Label
a	No	No
b	No	Yes
c	Yes	No
d	Yes	Yes

Table 2. The statistics of each metadata attribute within image categories a, b, c and d.

	Brandenburg Gate				**Reichstag**			
	a (30%)	**b (19%)**	**c (11%)**	**d (40%)**	**a (27%)**	**b (11%)**	**c (25%)**	**d (37%)**
Avg. user tag count	18	8	13	11	35	12	22	10
Avg. user photo count	19,087	3,852	18,354	5,422	8,136	7,928	9,555	2,618
Avg. user contact count	338	111	134	132	108	141	153	110
Avg. distance to the Target (m)	626.5	402.9	299.1	161.6	1,321	735.9	510.5	436.6

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, Figure 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 to 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 lower number of contacts in comparison to users falling in to the remaining categories. Hence, popularity in Flickr may not be a priority for this group of users, while priority in quality is.

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 category **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, 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 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 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.

Further, we have computed the distance to the target by taking the orthodrome between the geotag and the actual geographical coordinates of a point of interest. This reveals that the average distance to the target decreases for images from **a** to **d** within the use cases for Brandenburg Gate as well as Reichstag. Images in category **a** have the highest averaged distance to the target and in category **d** have the least averaged distance to the target (Table 2). The closer to the point of interest 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 point of interest, the person might become more imprecise when geotagging and labelling the image.

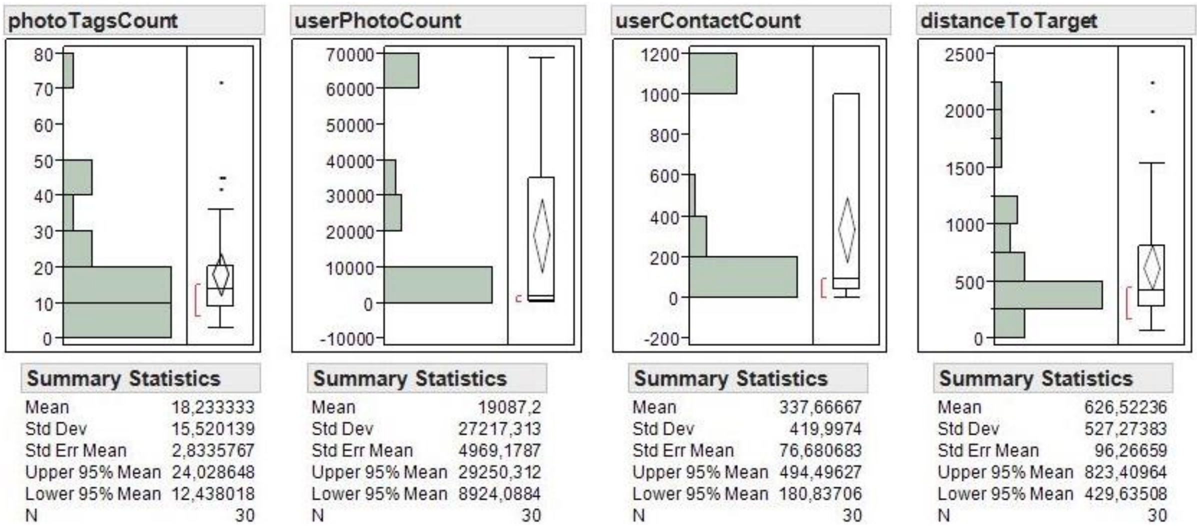

Figure 5. Distribution of data for category ‘a’ within the Brandenburg Gate use case.

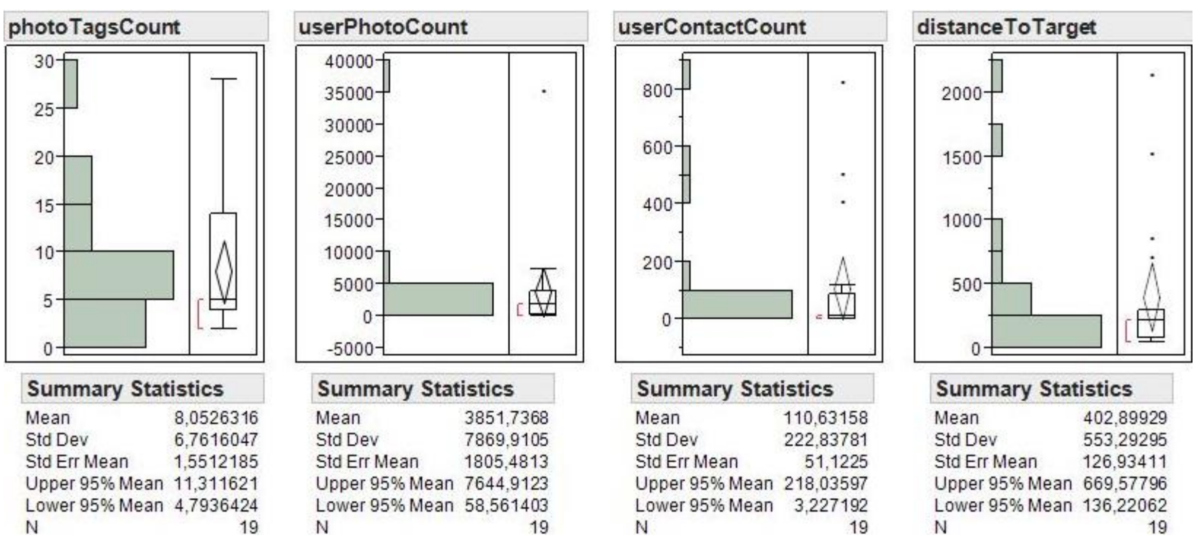

Figure 6. Distribution of data for category ‘b’ within the Brandenburg Gate use case.

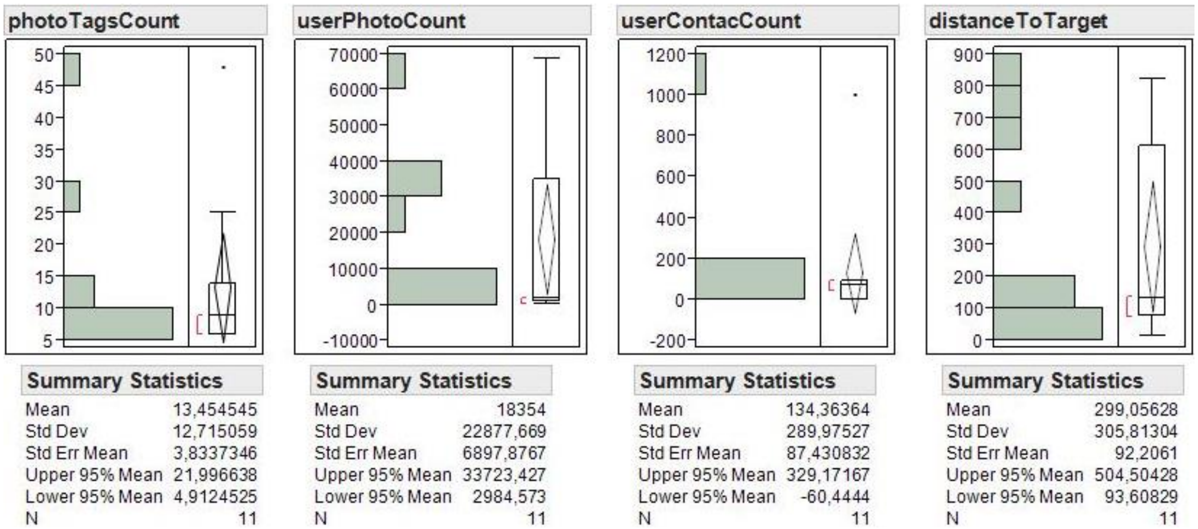

Figure 7. Distribution of data for category ‘c’ within the Brandenburg Gate use case.

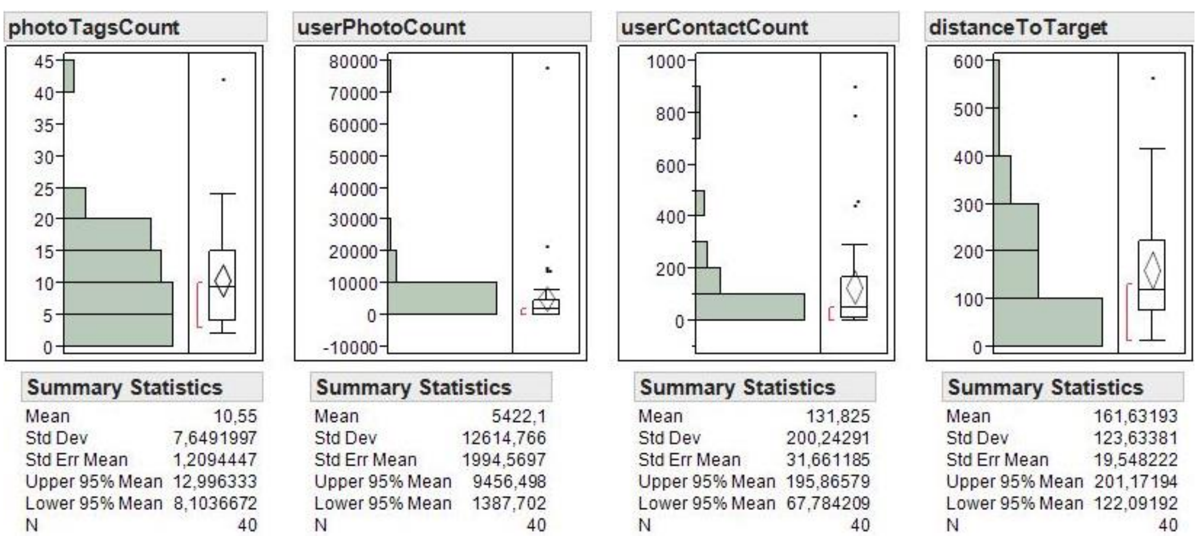

Figure 8. Distribution of data for category ‘d’ within the Brandenburg Gate use case.

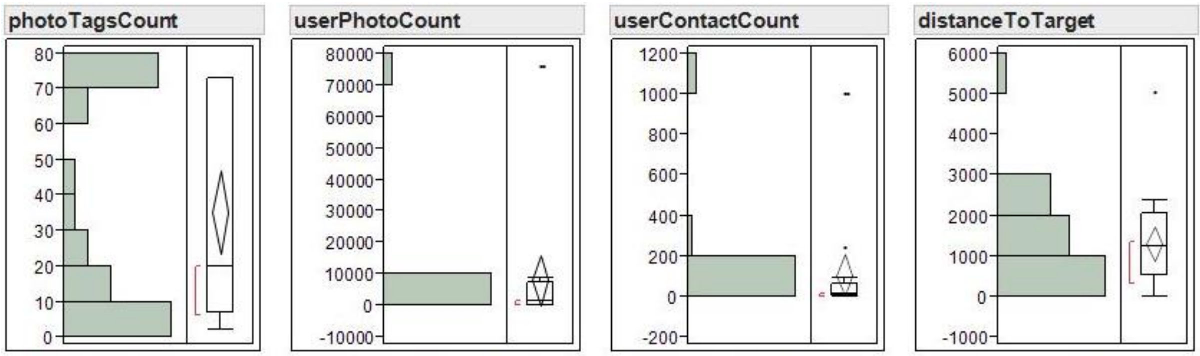

Summary Statistics	photoTagsCount	userPhotoCount	userContactCount	distanceToTarget
Mean	35,296296	8136,0741	107,92593	1320,6647
Std Dev	29,83134	19785,186	261,64773	1052,0642
Std Err Mean	5,741044	3807,6608	50,354128	202,46985
Upper 95% Mean	47,097181	15962,833	211,43032	1736,8474
Lower 95% Mean	23,495411	309,31522	4,4215326	904,48196
N	27	27	27	27

Figure 9. Distribution of data for category ‘a’ within the Reichstag use case.

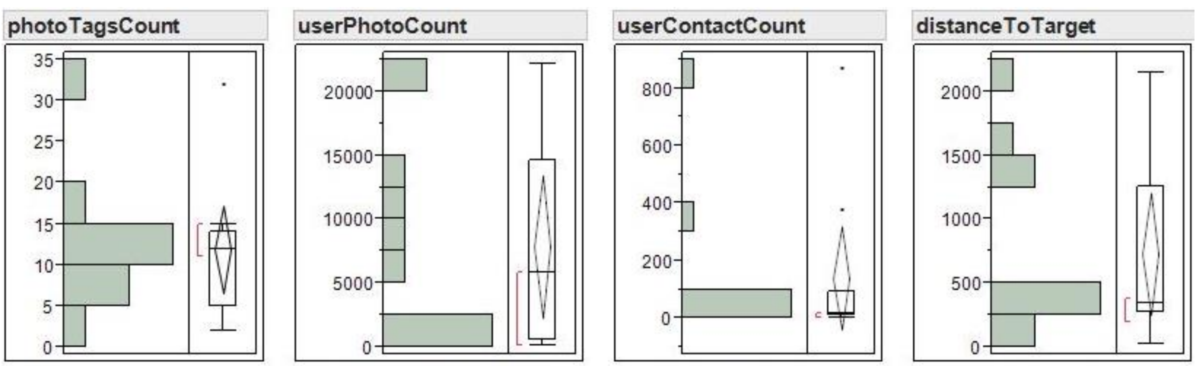

Summary Statistics	photoTagsCount	userPhotoCount	userContactCount	distanceToTarget
Mean	11,909091	7927,9091	141,18182	735,91635
Std Dev	7,892459	8287,3057	264,84215	712,84539
Std Err Mean	2,3796659	2498,7167	79,852912	214,93097
Upper 95% Mean	17,211317	13495,397	319,10519	1214,8124
Lower 95% Mean	6,6068648	2360,4213	-36,74156	257,0203
N	11	11	11	11

Figure 10. Distribution of data for category ‘b’ within the Reichstag use case.

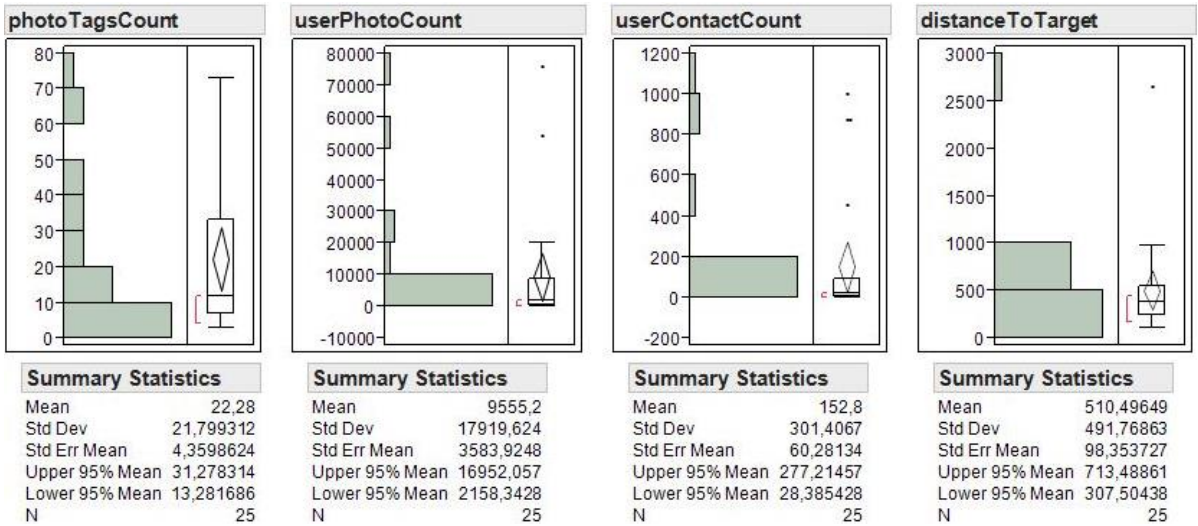

Figure 11. Distribution of data for category ‘c’ within the Reichstag use case.

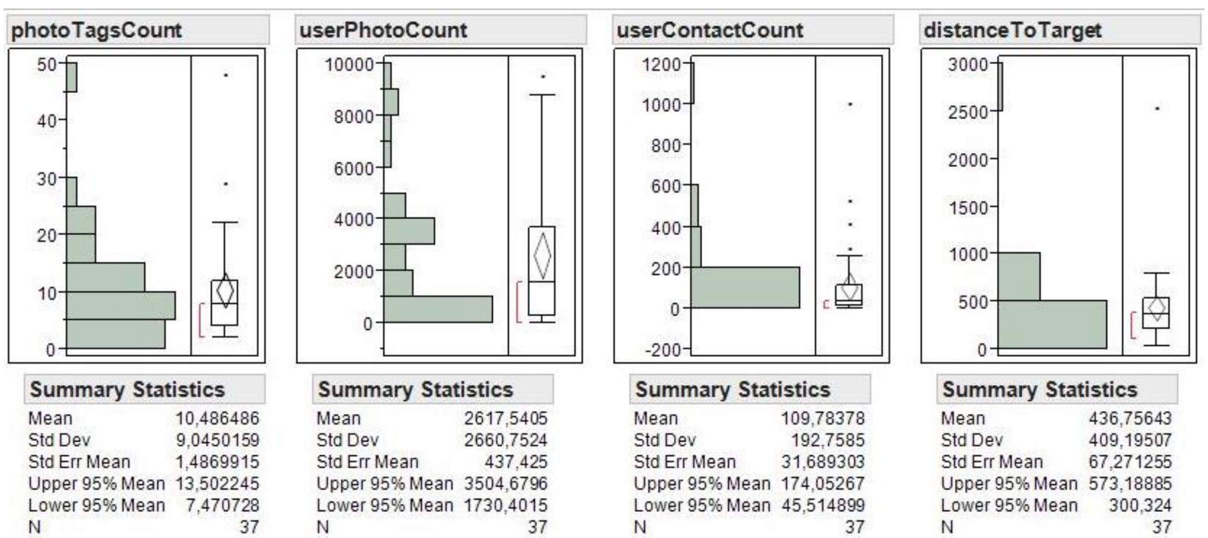

Figure 12. Distribution of data for category ‘d’ within the Reichstag use case.

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 correctness, other features such as the label precision, or image content can be used to evaluate the user credibility. Methods to utilise 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 point of interest through the geotag and the image label. Images placed within a visibility region that do not include the position of the point of interest are determined to be either incorrectly geotagged, or 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 mistakes can for example result from either manually adding wrongly measured coordinates as a geotag to the photo, or coordinates measured by a malfunctioning GPS device. Another reason might be that, while the geotag is correct, a user lacks sufficient knowledge about what is shown on the photograph and provides incorrect place describing tags. Also, 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 are photos of a certain place and a user draws comparisons to other sights by also 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 the Times Square in New York and provides according tags. Hence, a complement to our approach would be to utilise image recognition techniques that can programmatically identify the image content and compares it to the point of interest to find (dis)similarities, and then associate it with the reverse viewshed. This would already 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 aid us in filtering out relevant and irrelevantly 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 attribute 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 emphasised by Fisher (1991). Those quality issues are due to data errors, data resolution, as well as errors in the viewshed analysis algorithm. Thus, within this paper we limit our approach to calculating a reverse viewshed upon which the location correctness of geotagged Flickr images are assessed. We propose to use other additional user/photo metadata in combination with the location correctness to infer the credibility of users as an extension to future work.

6. CONCLUSIONS & FUTURE WORK

This paper contributes to the research and discussion on quality control of VGI. We have investigated through experimental analysis how a reverse viewshed analysis can be utilised to assess the location correctness of visually generated VGI. In doing so, we have first programmatically downloaded metadata of photographs for a certain point of interest by querying the open Flickr API for all geotagged photos, which 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 points of interest, 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 from a given observer point. If it lies outside of 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 point of interest through the geo/text tag do not necessarily visually represent the point of interest. This is also exhibited through analysing a sample dataset. We propose to conduct in the future work image recognition techniques to filter out images that are irrelevant to the point of interest.

Within the sample dataset for Brandenburg Gate and Reichstag we have categorised the photographs into four groups based on the geotag and label correctness. On those categories we made observations in 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). Also, we found that photos in category **a** and **c** (both incorrectly labelled) have higher numbers of tags. Further, the producers of photos in category b and d (correctly labelled) together have on average lower number of contacts as compared to 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 on investigating the tagging behaviour of users beyond their motivational aspects.

For the future, we will work towards a mechanism for automatically inferring the user credibility through analysing 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 vectorised city models based on CityGML. Further, we will look into the possibility of extracting credibility-related measures from analysing 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 behaviour. 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 to train statistical prediction algorithms for classifying the users according to the above mentioned weighted reliability parameters.

7. REFERENCES

1. Ames, M., Naaman, M. 2007. Why We Tag: Motivations for Annotation in Mobile and Online Media. In CHI '07: Proc. SIGCHI Conference on Human Factors in Computing Systems (San Jose, CA, 2007), ACM, 971–980.
2. Becker, C., Bizer, C. 2011. flickr™ wrappr- Precise Photo Association. Online available: http://wifo5-03.informatik.uni-mannheim.de/flickrwrappr/ (retrieved on 12.09.2012).
3. Bishr, M., Kuhn, W. 2007. Geospatial Information Bottom-Up: A Matter of Trust and Semantics. The European Information Society, 365-387.
4. Castillo, C., Mendoza, M., Poblete, B. 2011. Information Credibility on Twitter. In Proceedings of the 20th international conference on World Wide Web, ACM, 675-684.
5. 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(1), 332-358.
6. Ciepłuch, B., Jacob, R., Mooney, P., 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, July 20-23, 2010 (337). University of Leicester.
7. Crandall, D.J., Backstrom, L., Huttenlocher, D., Kleinberg, J. 2009. Mapping The World's Photos. In WWW '09: Proc. 18th international conference on World Wide Web (Madrid, Spain, 2009), ACM, 761–770.
8. De Floriani, L., Marzano, L., Puppo, P.E. 1994. Line-of-sight Communication on Terrain Models. International Journal of Geographic Information Systems, 8(4), 329-342.
9. Flanagin, A.J., Metzger. M.J. 2008. The Credibility of Volunteered Geographic Information. GeoJournal, 72(3), 137-148.
10. Friedland, G., Choi, J., Lei, H., Janin, A. 2011. Multimodal Location Estimation on Flickr Videos. ACM MULTIMEDIA, November 28 - December 1, 2011, Scottsdale, AZ, USA.
11. Frew, J. 2007. Provenance and Volunteered Geographic Information. Online available: http://www.ncgia.ucsb.edu/projects/vgi/docs/position/Frew_paper.pdf (retrieved on 10.07.2012).
12. Fisher, P.F. 1991. First Experiments in Viewshed Uncertainty: The Accuracy of the Viewshed Area. Photogrammetric Engineering and Remote Sensing 57, 1321-1327.
13. Fisher, P.F. 1993. Algorithm and Implementation Uncertainty in Viewshed Anlysis. International Journal of Geographic Information Systems 7(4), 331-347.
14. Fisher, P.F. 1996. Extending the Applicability of Viewsheds in Landscape Planning. Photogrammetric Engineering & Remote Sensing 62(11), 1297-1302.
15. Girres, J.F., Touya, G. 2010. Quality Assessment of the French OpenStreetMap Dataset. Transactions in GIS, 14(4), 435-459.
16. Golder, S., Huberman, B. 2006. Usage patterns of collaborative tagging systems. Journal of Information Science 32(2), 198-208.
17. Goodchild, M.F. 2007. Citizens as Sensors: The World of Volunteered Geography. GeoJournal, 69, 211-221.
18. Goodchild, M.F. 2009. NeoGeography and the Nature of Geographic Expertise. Journal of Location Based Services, 3(2), 82-96.
19. Goodchild, M.F., Li, L. 2012. Assuring the Quality of Volunteered Geographic Information. Spatial Statistics, 1, 110-120.
20. Golder, S.A., Huberman, B.A. 2006. Usage Patterns of Collaborative Tagging Systems. Journal of Information Science, 32(2), 198-208.
21. Girardin, F., Blat, J., Calabrese, F., Fiore, F.D., Ratti, C. 2008. Digital Footprinting: Uncovering Tourists with User-generated Content. IEEE Pervasive Computing 7(4), 36–43.
22. Gupta, M., Zhao, P., Han, J. 2012. Evaluating Event Credibility on Twitter. In SIAM International Conference on Data Mining (SDM 2012), April 26-28, 2012, Anaheim, CA, USA.

23. Haklay, M. 2010. How good is Volunteered Geographical Information? A Comparative Study of OpenStreetMap and Ordnance Survey Datasets. Environment and Planning. B: Planning & Design, 4(37), 682-703.
24. Haklay, M.M., Basiouka, S., Antoniou, V., Ather, A. 2010. How Many Volunteers Does it Take to Map an Area Well? The Validity of Linus' Law to Volunteered Geographic Information. The Cartographic Journal, 47(4), 315-322.
25. 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.
26. Hovland, C.I., Janis, I.L., Kelley, J.J. 1953. Communication and Persuasion. New Haven, CT, USA: Yale University Press.
27. Jankowski, P., Andrienko, N., Andrienko, G., Kisilevich, S. 2010. Discovering Landmark Preferences and Movement Patterns from Photo Postings. Transactions in GIS, 14(6), 833-852.
28. Kidner, D., Sparkes, A., Dorey, M. 1999. GIS and Wind Farm Planning. In Stilwill, J.,Geertman, S., Openshaw, S. (Eds), Geographical Information and Planning, Springer, London. 203-223.
29. Kim, Y., Rana, S., Wise, S. 2004. Exploring Multiple Viewshed Analysis Using Terrain Features and Optimisation Techniques. Computers and Geosciences, 30(9), 1019-1032.
30. Lake, I.R., Lovett, A.A., Bateman, I.J., Langford, I.H. 1998. Modelling Environmental Influences on Property Prices in an Urban Environment. Computers, Environment and Urban Systems, 22(2), 121-136.
31. Montello, D.R., Goodchild, M.F., Gottsegen, J., Fohl, P. 2003. Where's Downtown? Behavioral Methods for Determining Referents of Vague Spatial Queries. Spatial Cognition and Computation 3(2–3), 185–204.
32. Moxley, E., Kleban, J., Manjunath, B.S. 2008. Spirittagger: A Geo-Aware Tag Suggestion Tool Mined from Flickr. In Proceeding of the 1st ACM international conference on Multimedia Information Retrieval, ACM, 24-30.
33. O'Sullivan, D., Unwin, D.J. 2003. Geographic Information Analysis.
34. Ralling,P., Kidner, D., Ware, A. 1999. Distributed Viewshed Analysis for Planning Application. In Innovations in GIS 6. Getting, B. (Ed), Taylor and Francis, London, 185 - 199.
35. Rinner, C., Kessler, C., Andrulis, S. 2008. The Use of Web 2.0 Concepts to Support Deliberation in Spatial Decision-Making. Computers, Environment and Urban Systems, 32, 386-395.
36. Samet, H. 1984. The Quadtree and Related Hierarchical Data Structures. ACM Computer. Surveys, 16(2), 187-260.
37. Sherren, K., Fischer, J., Pink, J., Stott, J., Stein, J., Yoon, H. 2011. Australian Graziers Value Sparse Trees in Their Pastures: A Viewshed Analysis of Photo-Elicitation. Society & Natural Resources 24(4), 412-422.
38. Surowiecki, J., 2005. The Wisdom of Crowds. Anchor, New York.
39. Sigurbjoernsson, B., Van Zwol, R. 2008. Flickr Tag Recommendation based on Collective Knowledge. In WWW '08: Proc. 17th international conference on World Wide Web (Beijing, China, 2008), ACM, 327–336.
40. Valli, C., Hannay, P. (2010) Geotagging Where Cyberspace Comes to Your Place. Security and Management, 627-632.
41. Wang, J., Robinson, G.J., White, K. 1996. A Fast Solution to Local Viewshed Computation Using Grid-Based Digital Elevation Models. Photogrammetric Engineering & Remote Sensing, 62(10), 1157-1164.
Zwol, R.V. 2007. Flickr: Who is Looking? IEEE/WIC/ACM International Conference on Web Intelligence, Silicon Valley, USA (November 2-5, 2007).