

Linking VGI for place-based map generalization

Azam Bahrehdar
Department of Geography
Winterthurerstrasse 190
Zurich, Switzerland
azam.bahrehdar@geo.uzh.ch

Ross S. Purves
Department of Geography
Winterthurerstrasse 190
Zurich, Switzerland
ross.purves@geo.uzh.ch

Abstract

Traditional map production is a top-down process driven by national mapping agencies. We argue that VGI have the potential to allow us to create place-based maps, and in particular that VGI in the form of Flickr data can be used to drive generalization operators applied to Open Street Map data. We introduce a method to link the descriptions associated with georeferenced Flickr images to Open Street Map geometries for popular London locations, and demonstrate that rich semantic data, contributed by many users, are available after filtering for bias. We illustrate the potential application of our approach to three generalization operators.

Keywords: Place-based maps, map generalization, VGI

1 Introduction

Traditional map production is based on a top-down process, where the decisions about what to place on the map are driven by national mapping agencies (NMAs) and the legal and political framework within which they are positioned. In recent years, it has been argued that VGI, or more generally User Generated Content (UGC) offer the possibility to derive platial information – that is to say information that reflects notions of place as a lived and experienced space [1]. However, much of this research has concentrated on either extracting properties of space, for example in the form of vernacular placenames or place properties as an end to itself, and not as an additional form of semantics for use in the production of cartographic representations.

Parallel to the growth of VGI have been developments, which move maps away from being static products, and see them services that can adapt to individual users and use contexts. One obvious way of theoretically underpinning such adaptive maps, and the generalization processes necessary to produce them, is consideration of the notion of place as a basis for adaptation of content. If VGI contain platial information, then it should be possible to use this as an input to a holistic generalization process, both in making choices about which operators are important, and as parameters for individual algorithms.

In this paper we explore the potential of one rich source of VGI platial information, the tags associated with georeferenced Flickr images, in the process of generalizing a second VGI dataset in the form of OpenStreetMap (OSM). In order to use Flickr images to inform the generalization process we can consider three research questions, each of which we briefly explore in this paper:

RQ1: What forms of information can be extracted from VGI which can be exploited in the generalization process?

RQ2: How can information extracted from VGI in one form be linked to VGI in a second form?

RQ3: How can the information which has been extracted and linked be used explicitly in the generalization process?

The remainder of this paper is structured as follows. We initially briefly set out some key related work, firstly with respect to extracting information from Flickr, and secondly with respect to map generalization operations. We then introduce the data used in this pilot study, before explaining our methodological approach. Key results, illustrating how our methods could be used in a number of exemplar generalization operations are then presented, before we discuss our planned future work.

2 Related Work

The potential of Flickr as an information source for a wide range of geographic information, both in the form of geometry and semantics has long been recognised. Thus, for example, meaningful descriptions linked to location have been extracted by use of methods which privilege local over global information, and often are dominated by toponyms [2]. Other work has sought to extract the geometric footprints associated with both administrative and vernacular toponyms, thus linking placenames to a specific location [3]. The potential of Flickr to contain information related to both visible parts of a scene and their properties has been analysed by a number of authors [4]. Moreover, the spatial pattern of photographs itself has been recognised to provide information about the digital footprints of those visiting a location, indicating for example particularly commonly photographed locations [5][6]. Notwithstanding the great potential of VGI, any analysis of Flickr tags associated with images must also consider issues of bias and the underlying distribution of data [3][4].

OSM has also been the focus of a very wide range of research, initially with respect to classic GIScience themes including quality, reliability, and completeness with respect to both geometry and thematic information, and latterly also taking up topics such as gender bias or the use of OSM in routing and navigation [7][8][9]. However, here our focus is on the development of approaches for the generalization of OSM based on platial information, that is to say, methods “characterized by place names and descriptions as well as semantic relationships between places” [10]. To our knowledge, little or no research has been carried out on the automatic production of place-based maps, perhaps with the exception of ideas developed in computer graphics on the automatic generation of tourist maps [11].

In map generalization several very useful taxonomies of operators and conceptual models of the process have been developed which emphasise the importance of map purpose [12], [13]. Map purpose in turn implies that map semantics (i.e. what is being displayed) are central to the generalization process – however, much research has focused on geometric transformations for specific feature classes (i.e. the simplification of mountain roads while retaining essential features). We believe that VGI offers a potentially very rich source of semantic information which, by linking descriptions to existing geometries, can provide us with a novel way of steering the map generalization process holistically which go beyond the semantics captured in top-down, administratively generated, topographic data. However, in this paper we focus our attention on a preliminary study of the use of VGI in individual generalization operators.

3 Methodology

3.1 Description of data

In our experiment, we have restricted ourselves to ten popular places in London. The first four places are amongst the top seven photographed landmarks on Earth (such as Trafalgar Square, Tate Modern, Big Ben, and London Eye) and the next three places are in a list of top seven photographed landmarks in London (Piccadilly Circus, Buckingham Palace, Tower Bridge) [5]. The last three, St. Paul’s Cathedral, the Globe Theatre and Hyde Park, were chosen as prominent tourist attractions.

We extracted footprints for each location from OSM, and using these footprints and a freely available 1m Digital Surface Model (DSM) of London identified a nominal visible area associated with each geometry. Within the bounding box of London we retrieved, for all georeferenced Flickr images, tags, positions, user ids, taken and posted times, accuracy information, titles and descriptions using the Flickr API.

3.2 Removing Bias

As discussed above, raw VGI are prone to a range of biases. We firstly filtered out all images with an accuracy level of less than 15 (according to Flickr 11 corresponds to city level accuracy [sic] and 15 street level). We also removed all data contributed by users who had only uploaded a single image, those who had subsequently deleted their profile or images and used regular expressions to attempt to remove images described only by machine generated tags.

3.3 Linking Flickr images to OSM

We identified images which might contain relevant descriptions for our named locations in a two stage process. Firstly, we searched for place names in textual information describing the images (tag, descriptions and titles) allowing some tolerance for misspelling words using Levenshtein distance [14]. This is important as placenames such as Buckingham Palace are often slightly misspelled.

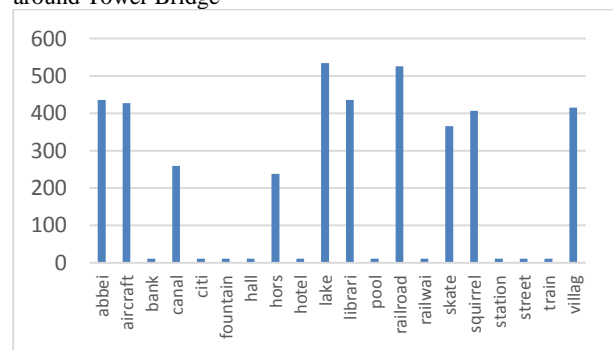
Secondly, we selected from these images only those which actually lay within the viewshed of our ten OSM geometries, thus including images which are more likely to be directly related to the object in question [15], [16].

3.4 Extraction of Place Descriptions

Having identified a set of images likely to be relevant to our ten named locations, we then attempted to extract information about the properties of these locations. To do so, we used a taxonomy based on visible elements in images (e.g. river, road, hill), qualities identifiable in images (e.g. summer, urban, sunset) and activities depicted (e.g. music, festival, birthday) which has previously proved to facilitate extraction of meaningful descriptions using a controlled list as a starting point [4]. We also filtered the tags for toponyms using a list generated using OSM within the bounding box of Greater London.

Frequent tags are not necessarily those most useful in characterizing a location, since they may be strongly influenced by contribution bias [3]. By generating so-called tag profiles, for individual locations, where use of tags is binned according to prolificness of users, we removed tags whose profile was associated with a high coefficient of variation (>200). High coefficients of variation are typically associated with contribution bias, as illustrated in Figure 1.

Figure 1: Coefficient of variation for 20 selected elements around Tower Bridge



Tags with a coefficient of variation over 200 include stems for abbey, aircraft, canal, railroad and village which all appear to be out of context in this setting. Terms such as bank, railway, street and train all have low coefficients of variation and are therefore retained to describe this location.

Having removed biased tags, we were left with, for each of our ten regions, a set of image locations associated with that place name, and image locations associated with elements, qualities and tags. Since places are effectively experienced as regions, we used a density-based clustering approach [17] to

identify sub-regions within a named place with similar characteristics, in this instance as captured by a single tag.

4 Results and Interpretation

In the following, we firstly present statistics illustrating the richness of the data describing our ten London landmarks. All locations except for the Globe Theatre are characterized by more than 1000 images, and with the exception of Piccadilly Circus, more than 1000 users contributed to these descriptions. After removing tags with high coefficients of variation, a rich set of tags capturing information about elements, qualities and activities remains.

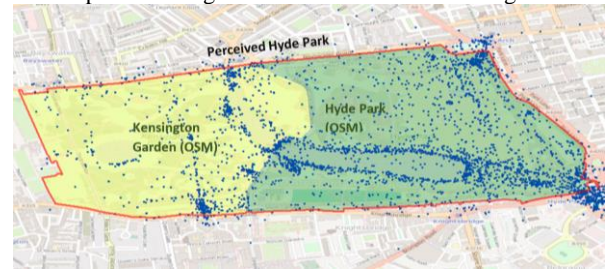
Table 1: Number of images, users and categorized tags for the ten locations after filtering and bias removal found in the visible area

Location	#images	#users	#elements, #qualities, #activities
Trafalgar Square	12801	3586	232,142,85
Tate Modern	7249	2490	227,135,74
Big Ben	8257	4335	229,134,72
London Eye	12645	4817	233,132,80
Piccadilly Circus	1102	691	124,88,49
Buckingham Palace	5733	2209	204,131,59
Tower Bridge	7738	3621	227,136,70
St. Paul's Cathedral	5430	2179	221,133,69
Globe Theatre	592	415	141,80,42
Hyde Park	6901	1761	229,133,76

For three places, we show the use of semantics derived from our Flickr data in generalization. These operations are all illustrative, and aim to demonstrate the potential of such data in the generalization process.

A common generalization operator is aggregation, where geometries with shared semantics are merged. In Figure 2 we show all points tagged with Hyde Park and the geometries of Hyde Park and the adjacent Kensington Gardens. It appears that Kensington Gardens is often perceived by visitors as being part of Hyde Park. Thus storing an aggregated geometry representing this perceived region, as well as the historical names more likely to be recognized by local residents may be a useful representation for a place-based map aimed at visitors to London.

Figure 2: OSM geometries and aggregation based on perceived region associated with Flickr images.



A second generalization operation is exaggeration, where a particular feature is made more visually prominent. Figure 3 demonstrates the potential of using clustered tags (in this case associated with river *bank*) with respect to Tower Bridge. Here, the south bank of the Thames appears to be clearly preferred as a view point, and a generalization operator focusing on a map of Tower Bridge might choose to widen the line representing the river bank in this region.

Figure 3: Cluster of images associated with tag river *bank* and its potential use in exaggerating the representation of the river bank to the south of the Thames.



In our final example, we illustrate the use of the related enlargement operator. Here, an existing geometry is scaled (rather than caricatured as for an exaggeration operator). In this case, the polygon representing the Tate Modern is enlarged (Figure 4) to capture its importance as captured both by the number of images and users associated with it (Table 1).

Figure 4: Enlarging the footprint of the Tate Modern as a function of its popularity.



5 Conclusions

Our aim in this paper was not to deliver generalized maps *per se*, but rather to illustrate the potential of linking VGI data in driving generalization operators which focus on not only geometry, but also semantics as contributed by large numbers of users. For ten popular locations in London, we showed that not only many photographs taken by a large user group existed, but also that descriptive tags controlled for bias could be extracted. Based on the locations of tags, their links to named places, the semantics of individual tagged images and popularity, we illustrated the application of three generalization operators. It is important to note that the key challenge in this process lies in integrating generalization operators in a holistic generalization process to build true place-based maps. This paper represents a first step in this process, and demonstrates that for maps linked to named places sufficiently rich and diverse data exist to use maps focused on individual locations.

Acknowledgements

We gratefully acknowledge funding through the Swiss National Science Foundation Project PlaceGen (149823).

References

- [1] T. Cresswell, "Place: A Short Introduction," *Prof. Geogr.*, vol. 58, no. February, p. 153, 2009.
- [2] T. Rattenbury and M. Naaman, "Methods for extracting place semantics from Flickr tags," *ACM Trans. Web*, vol. 3, no. 1, pp. 1–30, Jan. 2009.
- [3] L. Hollenstein and R. Purves, "Exploring place through user-generated content: Using Flickr to describe city cores," *J. Spat. Inf. Sci.*, vol. 1, no. 1, pp. 21–48, Jul. 2010.
- [4] R. S. Purves, A. J. Edwardes, and J. Wood, "Describing place through user generated content," *First Monday*, vol. 16, no. 9, pp. 1–17, 2011.
- [5] D. J. Crandall, L. Backstrom, D. Huttenlocher, and J. Kleinberg, "Mapping the world's photos," *Proc. 18th Int. Conf. World wide web - WWW '09*, p. 761, 2009.
- [6] F. Girardin, F. Calabrese, F. D. Fiore, C. Ratti, and J. Blat, "Digital Footprinting: Uncovering Tourists with User-Generated Content," *IEEE Pervasive Comput.*, vol. 7, no. 4, pp. 36–43, Oct. 2008.
- [7] M. Haklay, "How good is volunteered geographical information? A comparative study of OpenStreetMap and Ordnance Survey datasets," *Environ. Plan. B Plan. Des.*, vol. 37, no. 4, pp. 682–703, 2010.
- [8] T. Koukoletsos, M. Haklay, and C. Ellul, "Assessing Data Completeness of VGI through an Automated Matching Procedure for Linear Data," *Trans. GIS*, vol. 16, no. 4, pp. 477–498, Aug. 2012.
- [9] E. Nuhn and W. Reinhardt, "Generation of landmarks from 3D city models and OSM data," pp. 24–27, 2012.
- [10] S. Gao, K. Janowicz, G. Mckenzie, and L. Li, "Towards platial joins and buffers in place-based GIS," *Proc. 1st ACM SIGSPATIAL Int. Work. Comput. Model. place*, pp. 1–8, 2013.
- [11] F. Grabler, M. Agrawala, R. W. Sumner, and M. Pauly, "Automatic generation of tourist maps," *ACM Trans. Graph.*, vol. 27, no. 3, p. 1, 2008.
- [12] K. S. McMaster, R. B., & Shea, *Generalization in digital cartography*. Washington, DC: Association of American Geographers., 1992.
- [13] R. Weibel and G. Dutton, "Generalising spatial data and dealing with multiple representations," *Geogr. Inf. Syst.*, pp. 125–156, 1999.
- [14] V. I. Levenshtein, "Binary codes capable of correcting deletions, insertions, and reversals," *Soviet Physics Doklady*, vol. 10, no. 8. pp. 707–710, 1966.
- [15] H. Senaratne, A. Bröring, and T. Schreck, "Using Reverse Viewshed Analysis to Assess the Location Correctness of Visually Generated VGI," *Trans. GIS*, vol. 17, no. 3, pp. 369–386, Jun. 2013.
- [16] P. Fisher, "Extending the applicability of viewsheds in landscape planning," *Photogramm. Eng. Remote Sensing*, vol. 62, pp. 1297–1302, 1996.
- [17] A. B. Elsner, J. B., & Kara, *Hurricanes of the North Atlantic: Climate and soci*. Oxford University Press on Demand, 1999.