

Research Article

Discovering Landmark Preferences and Movement Patterns from Photo Postings

Piotr Jankowski

*Department of Geography
San Diego State University*

Natalia Andrienko

*Fraunhofer Institute – Intelligent
Analysis and Information Systems*

Gennady Andrienko

*Fraunhofer Institute – Intelligent
Analysis and Information Systems*

Slava Kisilevich

*Department of Computer and
Information Science
University of Konstanz*

Abstract

This article presents a geovisual analytics approach to discovering people's preferences for landmarks and movement patterns from photos posted on the Flickr website. The approach combines an exploratory spatio-temporal analysis of geographic coordinates and dates representing locations and time of taking photos with basic thematic information available through the Google Maps Web mapping service, and interpretation of the analyzed area. The article describes data aggregation and filtering techniques to reduce the size of the dataset and focuses on information addressing research questions. The results of analysis for the Seattle metropolitan area help to distinguish between sites that are occasionally popular among the photographers and can be considered as potential attractions from sites that are regularly visited and already known as city landmarks. The analysis of photographers' movements across the metropolitan area shows that most photographers' itineraries are short and highly localized.

1 Introduction

This article presents a study of the geovisual analytics approach to discovering preferences for urban landmarks and pertinent travel itineraries revealed through photographs posted on the photo sharing and social networking website Flickr. In addition to photo uploading, geotagging, and organizing functions, Flickr offers its users web services to embed images in blogs and social media. The issue of geotagged

Address for correspondence: Piotr Jankowski, Department of Geography, San Diego State University, 5500 Campanile Drive, San Diego, CA 92182-4493, USA. E-mail: piotr@geography.sdsu.edu

photos used as a proxy representation of people's interest in landmarks was introduced by others (Abbasi et al. 2009, Crandall et al. 2009). The goal of this study is twofold: (1) explore the potential of publically volunteered photos for providing information about people's activities in space and time, and (2) experiment with geovisual analytic techniques for extracting this information. In addition to a research motivation there may be a number of practical reasons for taking an interest in people's preferences for urban locations/landmarks and corresponding travel itineraries, such as for example, city promotion and advertising, public safety, tourism, and civic-minded activities. A traditional method of collecting information about such preferences beyond anecdotal evidence is a resident/visitor preference survey. This method, however, is expensive and time consuming. In this article we present an alternative approach, in which the revealed preferences for landmarks are discovered by processing and analyzing publically available databases of photos posted on Flickr website. The postings may be regarded as an expression of preference for a given location/landmark revealed by the photo authors. In addition to finding out about geographical preferences for landmarks we also analyze the itineraries of landmark photographers and their spatio-temporal pattern. In the article we explore three questions related to landmark preferences and the movement behavior of photographers:

1. Which locations outside the category of the most frequently visited tourist attractions gained the attention of Flickr photographers?
2. What were the likely reasons for attracting the attention of photographers?
3. What was the spatio-temporal pattern of photographer movement to visit the locations?

The fundamental challenge in answering the above questions is the nature of available data, which is voluminous (at the time of this writing there were 4 billion (4×10^9) photos uploaded on Flickr (<http://blog.flickr.net/en/2009/10/12/4000000000>); the number of geotagged photos is more difficult to establish – we managed to crawl 65 million photos, but this number does not include privately held photos – and semantically inconsistent (some photos have titles identifying the photographed objects while others do not). What makes the data interesting are their spatio-temporal characteristics: geographic coordinates of location where a photograph was taken and time of taking the photograph. In the remainder of the article we present a spatio-temporal analysis of Flickr photos, in which the information about photo location and time was combined with available photo titles and human interpretation of the photographed objects to discover new potential landmarks. We describe the analytical approach in section two. In the following sections we present the results of geovisual analytics for discovering landmark preferences (section three) and movement patterns of photographers (section four). We conclude the article with the summary of findings and future steps for extending the presented work.

2 Analysis Approach

An exploratory approach to visualization and analysis of spatio-temporal data follows a straightforward workflow comprised of three steps: overview first, zoom and filter, then details-on-demand (Shneiderman 1996). According to Keim (2005), these steps

might not be applicable to very large and complex datasets, such as photos posted on the Flickr website. Therefore, it may be necessary to begin the analysis of data while applying computational methods for reducing the size and complexity, including data aggregation, dimensionality reduction, or feature extraction. Furthermore, it may not be possible to explore thoroughly the whole dataset but instead focus only on what is important. Therefore, Keim (2005) suggests another workflow, called Visual Analytics Mantra, comprised of the following three steps: (1) analyze first – show the important, (2) zoom, filter and analyze further, (3) details on demand. We suggest the following procedure, corresponding with the Visual Analytics Mantra, for the analysis of a large set of photographs. In the first step, the spatio-temporal dataset is aggregated and summarized. This step of the workflow responds to the challenge of dealing with large datasets – a common trait of spatio-temporal data where there are simply too many records to depict each record directly – and gives the analyst a chance to spot interesting patterns. In the second step, the aggregates of the original data are analyzed using exploratory visual and analytic tools, which by virtue of being linked provide various data views in geographic, attribute, and temporal spaces. In the third step, interesting patterns revealed in the course of exploratory analysis are investigated in more detail by focusing on subsets of original data underlying the observed patterns. This step may involve additional data transformations and exploratory analysis augmented by human interpretation of the photograph content to help interpret the observed data patterns.

There are different types of spatio-temporal data, which can be defined in terms of three basic component sets (Peuquet 1994):

- Space (set of places), denoted here by S
- Time (set of moments or intervals, jointly called time units), denoted here by T , and
- Set of objects, denoted here by O

Elements of each set may be characterized by appropriate attributes. Types of spatio-temporal data may be defined according to the kind of changes they describe:

- *Spatially referenced time series* describe changes of attribute values of objects or places. This may be denoted as $O \times T \rightarrow A_1 \times A_2 \times \dots \times A_n$ (for each pair of object and time unit there are specific values of attributes A_1, A_2, \dots, A_n) or $S \times T \rightarrow \{A_1 \times A_2 \times \dots \times A_n\}$ (for each pair of place and time unit there are specific attribute values).
- *Spatial events data* describe existential changes (appearance and disappearance of objects). We define an event as an object with limited life time, i.e. existing during some time moment or interval. A spatial event has a position in space. Spatial event data may be denoted as $O \rightarrow T \times S \{A_1 \times A_2 \times \dots \times A_n\}$, where the component T defines the moment or interval of object existence, S defines the spatial position, and the optional components A_1, A_2, \dots, A_n specify values of object attributes.
- *Movement data* describe changes of spatial positions of objects. This may be denoted as $O \times T \rightarrow S \{A_1 \times A_2 \times \dots \times A_n\}$ (for each pair of object and time unit there is a specific place in S and, optionally, specific values of attributes A_1, A_2, \dots, A_n).
- *Spatial flows data* describe binary links between places and their changes over time. In particular, such data may represent collective movements of objects: for each pair of places and time unit there is a certain number of objects that moved from the first place to the second and, possibly, additional collective characteristics of the objects. This may be generally denoted as $S \times S \times T \rightarrow \{A_1 \times A_2 \times \dots \times A_n\}$.

Spatial event and movement data can be transformed into spatially referenced time series by means of spatial and temporal aggregation. There are also aggregation techniques that transform movement data into spatial flows. Thus, in our analysis, we first deal with spatial event data (each photo shot is an event). In the process of analysis, we apply spatial and temporal aggregation to the events and thereby obtain time series of attribute values characterizing places. Next, from the original event data, we derive movement data. For this purpose, the photos made by the same photographer are linked into a chronologically ordered sequence, which is interpreted as the trajectory of the photographer. Then, we transform the movement data into spatial flows data.

To accomplish spatio-temporal aggregation, the geographical space is divided into compartments and the time span of the data is divided into intervals. The events fitting into each compartment and time interval are collected and their statistics, in particular, the number of events are computed. Consequently, each compartment is characterized by a time-series of event counts. Given the analysis task, the tessellation of the territory should be sufficiently fine and result in a potentially large number of compartments. However, the individual inspection of each compartment may be too expensive, given their number. A reasonable approach is grouping the compartments by the similarity of their characteristics and further processing of groups instead of individual compartments. Before any further processing of grouped compartments can happen, though, one has to answer the question of what are the characteristics differentiating potentially interesting locations from uninteresting ones. In the analysis of Flickr posted photos we used the following criteria to help discriminate between potentially interesting and uninteresting locations.

1. Locations never visited or visited only by a few photographers are uninteresting.
2. Locations which were consistently visited by a relatively high number of photographers are uninteresting (these are major tourist places and as such of no interest to us).
3. Locations that were consistently visited by a relatively moderate number of photographers are potentially interesting (not crowded with tourists but attracting stable attention).
4. Locations that periodically or occasionally attracted unusually many photographers are interesting. Such locations may be attractive for organizing special activities for tourists and/or locals.

Suitable characteristics for the classification of locations may be obtained by describing the time series of photo counts using simple descriptive statistics such as maximum and quartiles or percentiles and using the following selection conditions:

- A low maximum value means that the place was never sufficiently interesting.
- A high value of the 1st quartile means that the place was visited by relatively many photographers.
- A low value of the third quartile together with a high maximum value means that the place attracted high attention in less than 25% of the time intervals. This may be a periodically or occasionally visited place (but not necessarily). The use of percentiles such as 90, 95, or 99 can help to find places that attracted high attention in less than 10%, 5%, or 1% of time intervals, respectively. To figure out whether the increase of interest in these places was periodic or occasional, the analyst may need to look at the place's time series of photo postings ("details on demand").

The aggregation of movement data proceeds as follows. The trajectories are divided into segments corresponding to the division of time into intervals. Then, for each pair of spatial compartments s_1 and s_2 and time interval t all trajectory segments that start in s_1 and end in s_2 are collected and counted. The count gives the magnitude of the flow from s_1 to s_2 during the interval t .

3 Analysis of Places

The initial dataset was collected by downloading the photo metadata from the Flickr site using the publicly available API. The downloading of metadata was performed using a similar approach to the one described in Crandall et al. (2009); first, a user id was obtained to download his/her photo metadata. Then, we downloaded all the user's contacts. To speed up the process of retrieving heterogeneous users, we retrieved all groups to which the user belonged. Using the group information we were able to retrieve all group members. The process was repeated for each user. We collected 64,975,609 entries from 2,617,271 users to the time of writing this article, which allows us to apply spatio-temporal analysis of photos related to virtually any region in the world represented by the downloaded photos.

Geo-referenced (geotagged) photos may introduce positional uncertainty. Flickr offers tools through its web site application that allows a photo to be associated with its corresponding location at various scales ranging from a city or region all the way down to street level. Depending on the choice of scale and the memory of photographer, the accuracy of photo georeferencing will vary. The accuracy may be much improved for more recent postings of photos taken with GPS equipped digital cameras or GPS integrated wireless phones. In some cases, coordinates could refer to the position of the photographer, while in others they could refer to the location of the object being photographed. This justifies a human analyst involvement, in addition to automated analysis, in exploring the data. A review of the collected photos revealed that some temporal information was incorrect. There were 6,229 photos with dates in an incorrect format. In addition, 50,076 photos had dates later than October 1, 2009 while the data were collected before October 2009. There were also 72,454 photographers who uploaded just a single photo; those entries were not suitable for the analysis of photographers' movement patterns including the investigation of flows and routes connecting different places.

In the remainder of this section we describe the analysis of a subset of this Flickr dataset comprised of 577,053 geo-referenced photos of locations in the Puget Sound metropolitan area, Washington State, USA, including the cities of Seattle, Bellevue, Kirkland, Redmond and the neighboring communities, made by 9,324 photographers between January 1, 2005 and August 31, 2009.

3.1 Spatiotemporal Aggregation

We started the analysis by aggregating the locations of photos into spatial clusters with the radius of 500 m. The length of the radius is a rough approximation of the length/width of a parcel accommodating a city landmark, such as a building, park, stadium, etc. (Andrienko and Andrienko 2010). We used the cluster centroids as the seeds for generating 2,930 Voronoi polygons covering the analysis area. Unlike a regular grid, the

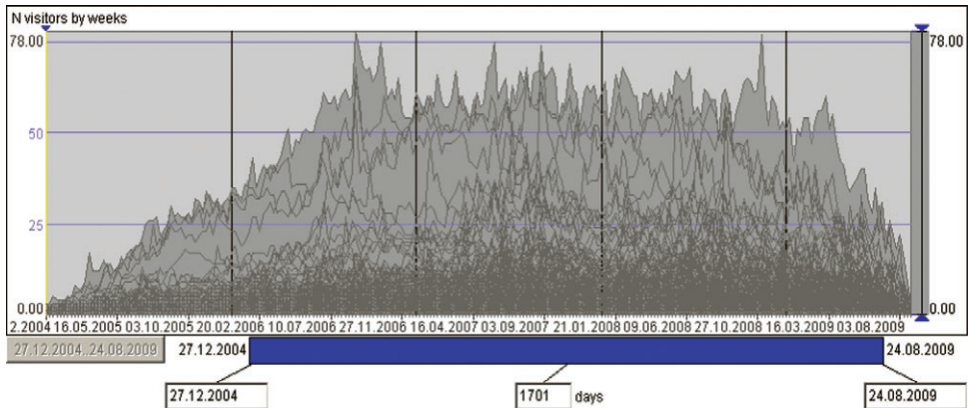


Figure 1 Each line in the graph represents the time-series of visits to a given place in the analysis area. There are 2,930 lines in total representing the corresponding number of Voronoi polygons

Voronoi polygon tessellation better matches the spatial distribution of photos, since the cluster centroids tend to be in the centers of spatial concentrations of photos that do not necessarily align with a regular shape of a grid cell. For each polygon, the total number of photos and the number of photographers were computed. The polygons are henceforth referred to in the article as “places”.

Next, we divided the 1/1/2005–31/8/2009 time period into 244 weekly intervals. For each place, the number of visits (number of photographers) in each week was computed. Hence, the data were transformed into the spatially referenced time series of 244 observations. The time series can be visualized on a time graph in Figure 1.

The time graph cannot be conveniently used for data exploration because of the large number of intersecting and overlapping lines. However, it is suitable for examining the time series of particular places by selecting only the corresponding time-series lines.

3.2 Interactive Grouping of the Places

For each place, the following statistics were computed from the 244 weekly observations: 1st quartile (25-percentile), median (50-percentile), 95-percentile, 99-percentile, and maximum. We used standard attribute query and classification tools to identify interesting places, characterized by high count values in the top percentiles. We also used the following classification corresponding to the criteria presented in section 2 to differentiate between potentially interesting (denoted by the asterisk symbol) and uninteresting places:

- Class 1: places with less than 10 visitors for the entire time-series (2,103 places) were classified as not interesting.
- Class 2: places with the maximum number of visitors in a week of less than five (548 places) were classified as not interesting.
- Class 3: places with the 1st quartile of weekly visitors greater than 10 (nine places) were regarded as uninteresting as these are the usual locations attracting many visitors otherwise known as the main tourist attractions.

Places with the maximum of weekly visitors ranging between 5 and 10 (213 places) were subdivided further into:

- Class 4*: places with relatively moderate values of the 1st quartile (25-percentile) and/or the median (50-percentile) indicating moderate but stable interest of photographers – these places (17) may be regarded as potentially interesting.
- Class 5: Places with relatively low values of the 1st quartile and/or the median signifying the lack of stable interest (196). Those places were deemed uninteresting.

For the remaining 68 places, we were interested in finding locations characterized by periodic or occasional peaks in the number of visitors. A high difference between the maximum number of visitors and the 95-percentile means that the place attracted high attention in less than 5% of the 244 time intervals (weeks), i.e. no more than in 12 weeks. A high difference between the maximum number of visitors and the 99-percentile means that the place attracted high attention in less than 1% of the time intervals (weeks), i.e. no more than in 2 weeks. To find places with such characteristics, we computed the differences between the maximum and the 95- and 99-percentiles. This resulted in two new attributes and their distribution values depicted in Figure 2. The scatterplot (Figure 2) reveals two interesting subsets of places: (1) points corresponding to high difference values between the maximum and both 95- and 99-percentiles (shown as dark hollow circles); (2) points corresponding to higher difference values between the maximum and 95-percentile, and lower difference values between the maximum and 99-percentile (shown as dark filled circles). Thereby one can create three additional classes:

- Class 6*: places with one or two occasional peaks of interest (five places represented by hollow points).




Figure 2 The distribution of differences between the maximum number of visitors and the 95- and 99-percentiles for 68 places in the Seattle metropolitan area

- Class 7*: places with more than two peaks of interest (six places represented by filled circles; there are two overlapping points in Figure 2), which, possibly occur periodically.
- Class 8*: the residual of 57 places characterized by the lack of occasional peaks representative of isolated one-week long intervals and by longer (2-3 consecutive weeks) periods of relatively high interest. These places are characterized by relatively high maximum values and low differences between the maximum and 95-percentile indicating locations that received high interest in more than 5% of time intervals, i.e. in more than 12 weeks. Additionally, these places are characterized by high differences between the maximum and 90-percentile thus indicating the number of “high interest” time intervals to be less than 10% or 24 weeks.

Next, we explored each place belonging to one of two classes in detail by visualizing their locations (Figure 3). The map reveals a spatial cluster formed by six places belonging to classes 6* and 7* in the Seattle’s Fremont district (Figure 3, location #1). Two of the places belong to class 6* representing high difference values between the maximum and both 95- and 99-percentiles (a few occasional peaks) while the other four belong to class 7* representing higher difference values between the maximum and 95-percentile, and lower difference values between the maximum and 99-percentile (still occasional but more frequent peaks than in class 6*). The other places, belonging to both classes and numbered from 2 to 6, are clustered around the Seattle downtown area.

The time series of places belonging to classes 6* and 7* and located in Seattle’s Fremont district are depicted in a time graph (Figure 4). The two highest numbers of visitor photographers in the time series happened in the summers of 2007 and 2008, more precisely, in the weeks starting from 11.06.2007 and 16.06.2008 and coincided with the Summer Solstice Parade & Pageant in Fremont – the artsy district of Seattle.

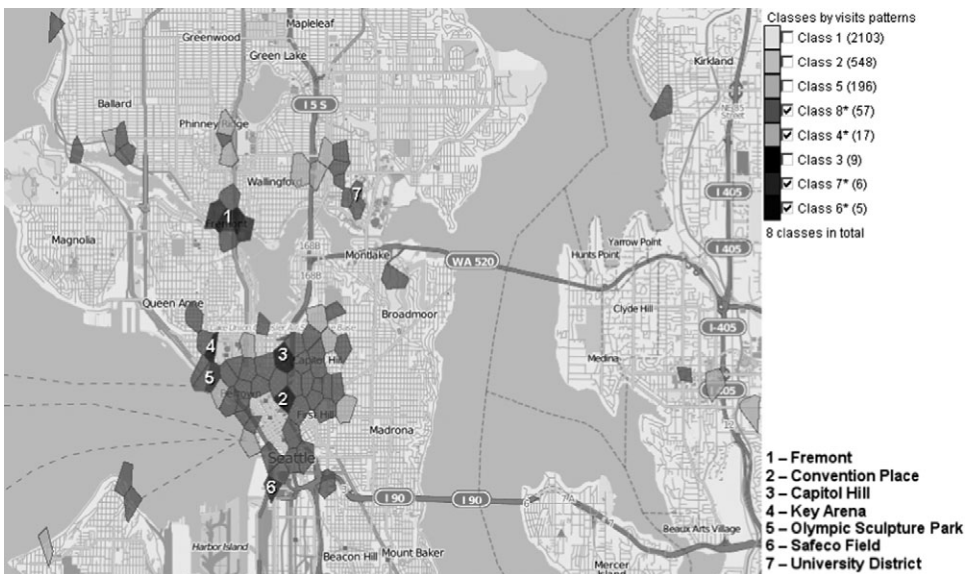


Figure 3 Grayscale shading indicates the places belonging to four selected classes (4*,6*,7*,8*) categorized earlier in the text as (potentially) interesting

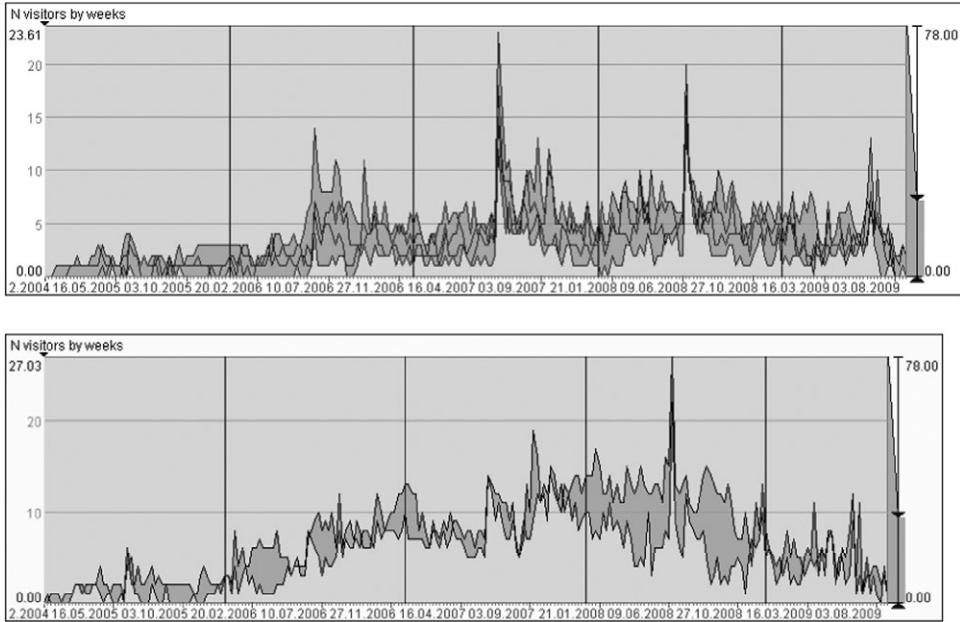


Figure 4 Time series of six places comprising cluster 1 on the map of study area (Figure 3) in the Fremont district of Seattle. The top part of the figure represents the time series for class 6* “high diff max – 95% AND high diff max – 99%” class while the lower part represents the time series for class 7* “higher diff max – 95% AND lower diff max – 99%”

Citing the http://en.wikipedia.org/wiki/Fremont,_Seattle,_Washington: “The Fremont Arts Council sponsors several highly attended annual events in Fremont. One of those events is the Summer Solstice Parade & Pageant, which has made Fremont famous for its nude Solstice Cyclists” (http://en.wikipedia.org/wiki/Solstice_Cyclists).

The interpretation of the cluster was corroborated by the inspection of corresponding photos. We retrieved the titles of the photos made in the places belonging to the cluster during the weeks with the highest number of visitors. Many titles indeed included “Fremont solstice parade”. In 2007, most of the photos were made on 16.06.2007, which is the Saturday before the summer solstice. In 2008, the parade took place on Saturday 21.06.2008. Searching the Internet sources we found out that in 2009 the parade was held on Saturday 20.06.2009. The corresponding peak in the time-series graph (Figure 4), as well as the peak of 12.06.2006 and a smaller one of 13.06.2005 point out to a periodically high interest associated with this regular event.

Further examination of the time-series for the cluster of places in the Fremont area revealed also a high number of visits re-occurring periodically in September, more precisely, in the weeks starting on 18.09.2006, 17.09.2007, and 15.09.2008, respectively. The inspection of the titles of photos provided an explanation for the higher than usual number of visits: the Fremont beer festival, which regularly takes place in the middle of September.

The remaining five places belonging to classes 6* and 7* and labeled 2–6 in Figure 3 do not form spatial clusters but they do concentrate around Seattle’s city center. The examination of their respective time series corroborated by the inspection of the photos

and related Internet sources revealed the following locations and reasons that attracted the attention of photographers:

- Place #2: Convention Place. Highest number of visitors: the weeks starting on 20.08.2007 and 25.08.2008. Explanation: game festivals PAX 2007 and PAX 2008 taking place at the Convention Place during 24.08–26.08.2007 and 29.08–31.08.2008.
- Place #3: Capitol Hill. Highest number of visitors: 15.12.2008. Explanation: snowfall and traffic accident.
- Place #4: Key Arena. Highest number of visitors: 04.02.2008. Explanation: Barack Obama's visit on 08.02.2008.
- Place #5: Olympic Sculpture Park. Highest number of visitors: the week of 15.01.2007–21.01.2007. Explanation: the Olympic Sculpture Park opened on January 20, 2007.
- Place #6: Safeco Field. Highest number of visitors: the time period from the week of 11–17.06.2006 until the week of 30.07–05.08.2007. Explanation: popular game season coinciding with better than average performance of the Seattle Mariners baseball team.

Further search for places of interest to photographers led us to examine the temporal patterns of photo taking activities in class 8* (Figure 3). The time series graph (Figure 1) can be transformed to show only line segments corresponding to substantial increases (or decreases) of the weekly number of visitors with respect to the previous week. Thus only the line segments representing the number of visitors at least twice as high as for the previous time interval (week) are visible in the time graph depicted in Figure 5. The switching between the standard mode and the line segmentation mode is accomplished by selecting and unselecting the checkbox “Show only selected line segments” at the bottom of the time graph.

In the segmentation mode (Figure 5), the analyst may select a particular line segment with the mouse (Figure 6) and then return the graph to the standard mode in order to see the whole time series line, to which the given segment belongs. The line is highlighted in black, as shown in Figure 6.

The examination of the time series characterized by repeatedly large increases in the weekly number of visitors led us to discover several unusual temporal patterns. A particularly interesting pattern, depicted in Figure 6, corresponds to location 7 on the




Figure 5 Time series graph shows only the line segments where the number of visitors was at least twice as high as in the previous interval




Figure 6 The time series of the line segment selected in the line segmentation mode as shown in Figure 5

map in Figure 3. The location is on the campus of University of Washington. The location attracted periodically a relatively high number of visitors during the weeks of 18–31.03.2006; 12–25.03.2007; 17–31.03.2008; 26.03–06.04.2009 (see the peaks in the time series in Figure 6). Since the increases are not occasional but consist of several high values during consecutive weeks (compare “wide” peaks in Figure 6 with “narrow” peaks in Figure 4), this pattern is representative of class 8*. The main campus of the University of Washington located in Seattle’s University District is famous for its early spring cherry tree blossoms attracting the local as well as national and international visitors.

4 Analysis of Photographers’ Movement in Space and Time

The objective of movement analysis was to find out whether there was a spatio-temporal pattern of visiting the locations of interest in the greater metropolitan area of Seattle. To facilitate the analysis the sequences of photos were extracted for each photographer who took more than one photo. A photo was regarded to be a part of a sequence if the time interval separating it from the previous photo was less than or equal to eight hours. Otherwise the photo was considered to be the beginning of a new photo session and hence a new sequence. This pre-processing of the data resulted in extracting 78,871 sequences created by 9,324 photographers. Thanks to their geographic coordinates photos belonging to a sequence can be plotted as point locations and aligned into a movement trajectory of photographers.

4.1 Spatial Analysis of Movement Trajectories

4.1.1 Aggregation

Movement trajectories can either be enclosed within one place (polygon container) or originate in one place and end in another. For every pair of places (A, B) the total number of times a photographer moved from A to B was counted. The resulting records (place A, place B, count) are called “aggregate moves”. There are also records of the type (place A, place A, count), which aggregate the trajectories fully contained in place A.

4.1.2 Visualization

Flow mapping (Slocum et al. 2009) is a standard cartographic technique to visualize aggregate moves. Since movement trajectories between any two places can proceed in

both directions, our movement visualization tool employs “half-arrow” symbols to represent movements between the places in two opposite directions. This symbol was proposed by Tobler (1987) for discrete flow maps. The symbols are integrated with lines connecting the centroids of the polygon compartments. The widths of the symbols are proportional to the numbers of aggregate moves they represent. The aggregate moves with coinciding starts and ends are represented by circular symbols with the radius proportional to the number of moves.

4.1.3 Results

1. More than one third of the trajectories (29,072 or 37%) were contained within one compartment. Such trajectories were transformed into the aggregate moves of the A, A, count type, i.e. where the start coincided with the end. Most of such moves were located in the city center. The circular symbols dominated the map and made the linear symbols representing the aggregate moves of the A, B, count type hardly visible. Hence, in the subsequent analysis the aggregate moves with the coinciding starts and ends were filtered out.
2. The remaining 49,799 trajectories represent the aggregate moves of the A, B, count type, i.e. where the end differs from the start. The large number of trajectories makes a visual pattern assessment nearly impossible and necessitates data filtering. We interactively filtered the trajectories by the place of start and/or end, the move count, and the length of the move (the distance between the centroids of place A and place B).
 - 2.1 Figure 7 depicts the trajectories filtered by the places of start and end, where the place location is restricted to the city center and the move count is at least 50. There are two distinct areas: Pacific Science Center (north-western section of the maps) and Pike Place Market (south-eastern section) where the places are connected with strong bi-directional flows. Moves in the south-central section are more frequent than in other sections of the city center.
 - 2.2 Figure 8 depicts the trajectories filtered by the places of start and end, where the start locations are in the city center and the end locations are outside the city center. Further filtering using the minimum number of moves along a trajectory equal to 10 reveals the movement pattern, in which the dominating trajectories connect the city center with the Green Lake and Magnolia areas (north and north-west of the city center), south-western Seattle, and the areas east of Seattle including the cities of Bellevue and Redmond. Filtering the trajectories using the reversed locations of start and end (starting outside the city center and ending in the center) returns a similar pattern.
 - 2.3 Figure 9 depicts the trajectories filtered by the locations of start and end being outside the city center and by the length of moves. The flow map on the left shows that the majority of trajectories, located outside the city center, are dominated by short moves (e.g. Fremont area and around). Long moves are much less frequent than the short moves; the lines depicting long moves are mostly thin representing a small number of photographers who moved over longer distances within a course of one photo

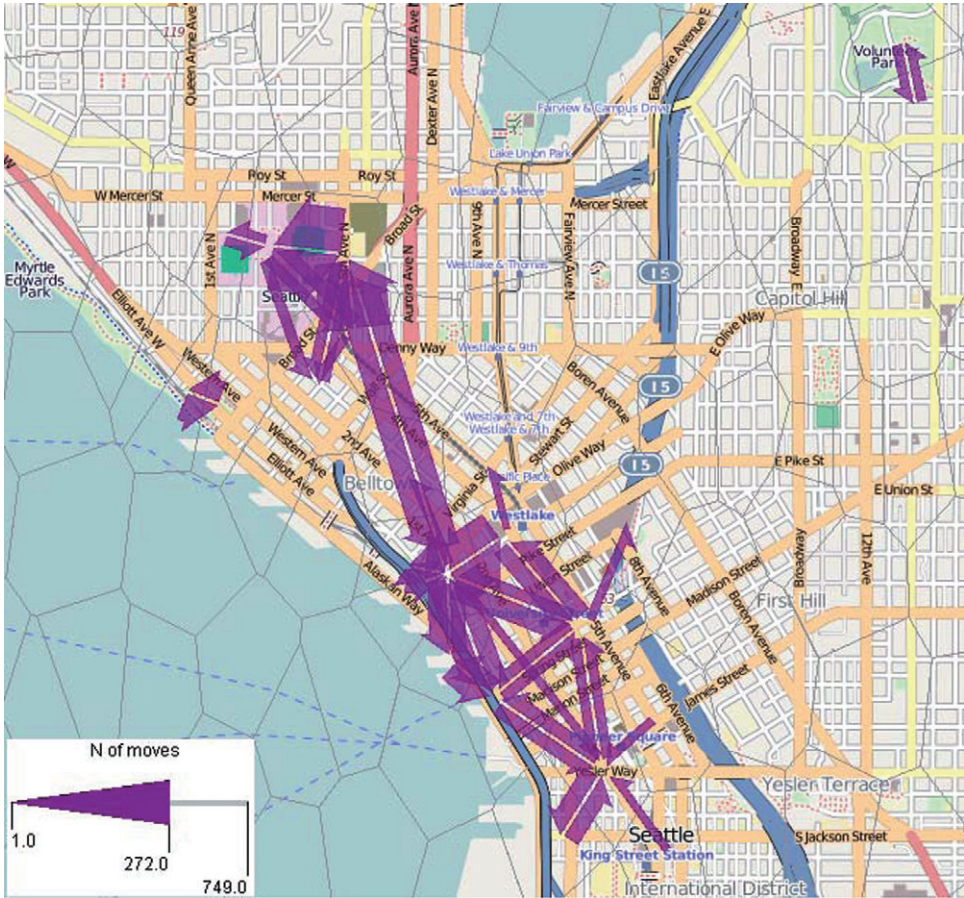


Figure 7 Trajectories in the center of Seattle filtered by the city center-only location of start and end, and by the minimum number of moves equal to 50. The irregular mesh of Voronoi polygons, representing places (compartments) is overlaid on the flow map

session. The largest move count for the short trajectories is 63 while for the long trajectories it is 20 (Figure 9). The flow map on the right shows that long trajectories most notably connected Bellevue with Redmond and Bellevue with Kirkland – cities on the east side of Lake Washington – and that there were only a few, low-count trajectories connecting different parts of the greater Seattle metropolitan area. The pattern of short moves *versus* long moves is consistent with the density of development in Seattle where the areas outside the city center are characterized by a fairly dense development in contrast to other neighboring cities (Bellevue, Kirkland, Redmond) where a low density, suburban development pattern is dominant outside the urban core. Only a small number of photographers moved over longer distances within a short time span of 8 hours. Most tended to visit fairly localized areas, as demonstrated by the predominance of short itineraries in Figure 9.




Figure 8 Trajectories originating in the center and ending outside the city center. The map on the left depicts all trajectories. The map on the right depicts only the trajectories with the move count of 10 or more

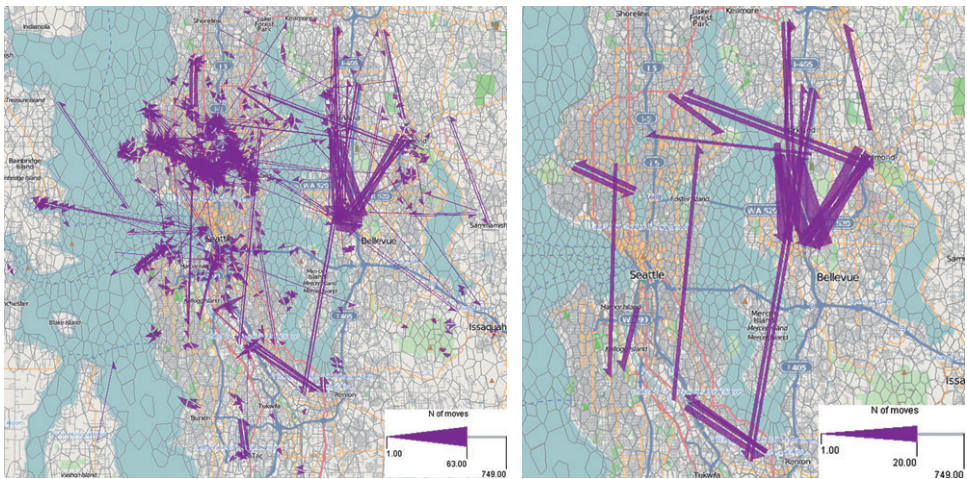


Figure 9 Trajectories originating and ending outside the city center. The flow map on the left shows the overall pattern comprised of short and long trajectories. The flow map on the right depicts the trajectories that are at the minimum 5 km long and have at least five moves

4.2 Spatio-Temporal Analysis of Movement Trajectories

The analysis was facilitated by data aggregation, in which for every pair of places (A, B) and time interval $[t_i, t_{i+1}]$, where i initially equaled one month, the number of moves from A to B was counted. For the monthly time interval there is no obvious trajectory pattern. We considered separately long moves (length ≥ 3 km) and short moves ($0 < \text{length}$

$\leq 1,500$ m). For long moves there was an absence of prominent re-occurring patterns that could be detected except for the trajectories across the Puget Sound, most notably between Seattle and Bainbridge Island, occurring mostly during summer months when weather is the best (Figure 10).

For the short moves ($0 < \text{length} \leq 1,500$ m), interesting patterns appear in some months. Some of them, such as those depicted in Figure 11, represent long sequences of photographing urban and suburban landscape, during which photographers moved and frequently took photos along the way in almost every compartment of the sequence. Such sequences might be analyzed further for their potential of becoming designated as possible scenic walking routes.

Next, the data aggregation was repeated for a longer time interval equal to three months and spatial clustering with a radius of 5,000 m instead of the initial radius of 500 m. The increased time aggregation interval reflects the seasonality of travel patterns




Figure 10 Long move trajectories (3 km or longer) for the month of June 2006. Similar pattern with moves from Seattle to Bainbridge Island (west) and back was observed for the summer months during the entire 2005–2009 period

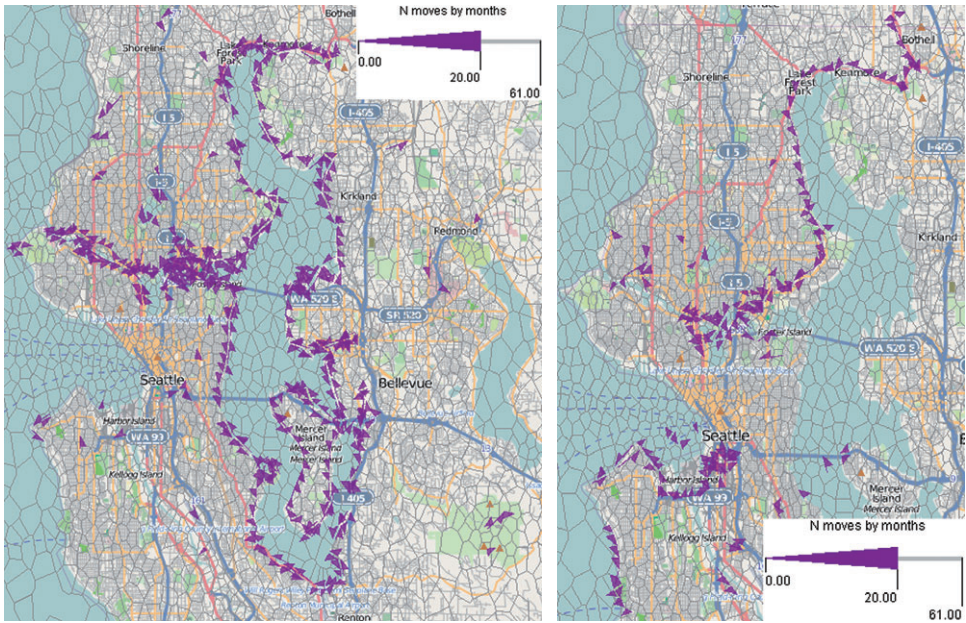


Figure 11 Short move trajectories ($0 < \text{length} \leq 1500$ m). The flow map on the left represents the aggregate moves for 02/2007. The map on the right represents the aggregate moves for 08/2009

in the study area, which roughly corresponds to quarters of the year. The increased radius of spatial clustering corresponds to our interest in capturing major flows in the study area without focusing necessarily on moves between specific landmarks. The results of analyzing movement patterns in space and time using the increased temporal and spatial clustering parameters are presented in Figure 12. Two patterns become apparent when summer trajectories are compared with winter trajectories across four years (2005–2008). First, the frequency of moves, represented by the width of the line, steadily increases across both seasons (summer and winter) from 2005 through 2008. This increase may simply reflect the fact of the growing popularity of Flickr resulting in the growing number of photographers between 2005 and 2008 posting their photos. Secondly, the seasonality (summer *versus* winter) is reflected only marginally in the movement trajectories connecting Seattle's downtown area with Bainbridge Island to the west of downtown and across Puget Sound. These trajectories are of low frequency representing a relatively low number of photographers moving across Puget Sound. The vast majority of trajectories concentrate, regardless of the season, around the city center and north of it. This pattern underscores the fact that Seattle has been a year-around destination for photographers.

4.3 Summary of Findings

The majority of photos analyzed in this study were taken within small areas (approximated by circles with radii below 500 m). This indicates that the photographers focused on specific locations or events rather than on photographing an extended metropolitan

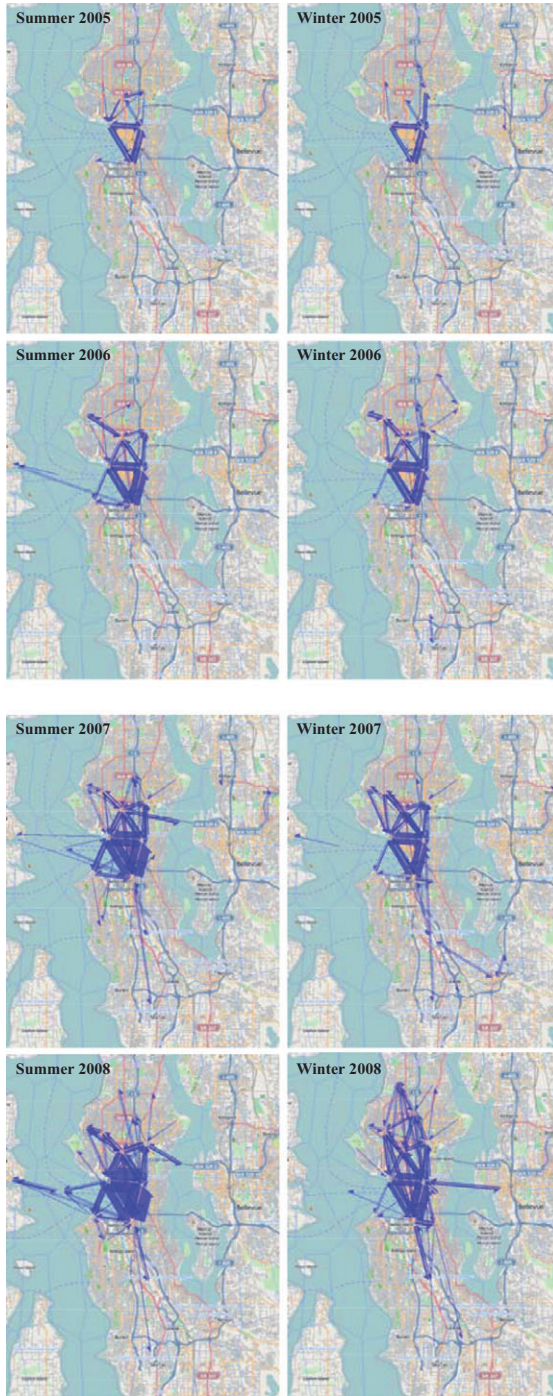


Figure 12 Pattern of movement trajectories in the study area aggregated by 3-month time interval and 5,000 m spatial cluster radius. The minimum line thickness (1 pixel) corresponds to five moves and the maximum thickness (36 pixels) corresponds to 176 moves. Flows with less than five moves are not represented

area. There were two larger areas in the city center accounting for most of the intra- and inter-area moves associated with photo taking activity. One area covers a large part of Seattle's downtown, from Pioneer Square to Pike Place Market. The other area extends around the Pacific Science Center north of Seattle's downtown.

Photographers' moves away from the city center to other locations were not as frequent as within the city center. Some of those locations to the north, south, and east (e.g. Redmond and Bellevue) of the city center are better connected than others by bi-directional flows of photographers. This means that during the period of 2005–2009 there were people interested in visiting in the same day the city center and some of outlying locations in the Puget Sound Metropolitan Area.

Both in the city center and outside of it, short moves between neighboring locations prevailed over distant moves. This may reflect a relative mobility of photographers who tour an area rather than staying in one location and often find interesting landmarks/events for taking photos as they move from one place to another. Long moves correspond to cases, in which photographers travel from one place to another without taking photos on the way. Such moves were relatively infrequent in our study. Some photographers were interested in walking or driving longer distances along Puget Sound, Lake Washington, and Lake Union waterfronts and taking photos on the way. There was a lack of clear seasonal dependency in the frequency and spatial direction of photographers' moves, meaning that similar move patterns occurred in different seasons (summer and winter).

5 Conclusions

The goals of this study were: (1) to explore the potential of volunteered geographic information, using the example of Flickr posted photos, for providing information about people's activities in space and time, and (2) to experiment with geovisual analytics techniques for extracting this information. In respect to the first goal, social networking websites such as Flickr and Panoramio, hosting databases of georeferenced photos, offer information on the:

- Spatiality of people's interests; locations of landmarks and events that are of interest to photographers,
- Temporality of people's interests; dates of photographing places and events and the seasonality of people's interests,
- Spatial extent of people's interests; boundaries of areas and events represented on photographs,
- Connectivity between photographed places represented by a network of moves connecting places of interest,
- Travel patterns of photographers and their temporal characteristics.

Similar information can be found in the databases of other social networking services storing spatial and temporal references of information created by their user, such as for example the georeferenced Twitter messages. Users of these services act as voluntary or sometimes involuntary sensors, collecting potentially useful geographic information (Goodchild 2007). At the same time, it is important to be aware of the limitations of volunteered geographic information including spatial and temporal coverage as well as demographic and social representativeness. The question of how closely do landmark

preferences and travel itineraries of social media users represent the preferences and itineraries of other groups, who do not use social media to communicate their preferences, remains unanswered by this study.

In respect to the second goal, the techniques of geovisual analytics used in this study proved to be effective in data aggregation and search for spatio-temporal patterns. There were two time-consuming tasks in the analysis of places (section 3) involving: (1) search for interesting temporal patterns of place visits, and (2) acquisition of additional information for interpreting the detected patterns (interpreting specific locations within their larger geographic context, retrieving and reading the titles of the photos, and searching the Web for information related to locations of interest). These tasks require an improvement in analytical support. More specifically, statistical techniques can help in detecting particular types of patterns such as sudden peak and periodic variation (Andrienko et al. 2010a). The interpretation of patterns could be supported in future studies by:

- Automated comparison of geotagged photo locations with locations of known objects (e.g. landmarks) stored in a geodatabase.
- Automated text analysis of photo titles and supplementary comments to extract place names and information identifying photographed events. Keywords extracted from photo titles and accompanying time references can be passed to a Web search machine for retrieving relevant information from the Web.

In the analysis of movement (section 4), the main problem of visual display clutter was overcome by interactive dynamic filtering, which enabled the visualization of filtered data. Further automation could be achieved by applying data mining techniques for extracting frequent item combinations and frequent item sequences (Kisilevich et al. 2010). In this case, items are the places occurring in the trajectories.

Further on, the investigation of the temporal variation of movement trajectories (section 4.2) could be supported by two-way clustering similar to the approach suggested by Andrienko et al. (2010b) for data aggregated by spatial compartments. In this approach, compartments are grouped according to the temporal variation of the respective attribute values (e.g. frequency of photographs). Complementary to this, time intervals are grouped according to the respective spatial distributions of the attribute values. The same idea can be applied to aggregated movement data by taking aggregate moves between places instead of spatial compartments. The time-variant counts of people are the attribute of the moves used for the clustering.

Acknowledgements

The first author of the paper gratefully acknowledges partial support for this research provided by the Science Foundation of Ireland ETS Walton Visiting Professorship. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of Science Foundation Ireland. The authors also wish to thank the reviewers for their valuable comments.

References

- Abbasi R, Chernov S, Nejd W, Paju R, and Staab S 2009 Exploiting Flickr tags and groups for finding landmark photos. In Boughanem M, Berrut C, Mothe J, and Soule-Dupuy C (eds)

- Advances in Information Retrieval: Proceedings of the Thirty-First European Conference on IR Research, ECIR 2009, Toulouse, France*. Berlin, Springer Verlag Lecture Notes in Computer Science Vol. 5478: 654–61
- Andrienko N and Andrienko G 2010 Spatial generalization and aggregation of massive movement data. *IEEE Transactions on Visualization and Computer Graphics* 16: in press
- Andrienko G, Andrienko N, Mladenov M, Mock M, and Pölitz C 2010a Discovering bits of place histories from people's activity traces. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology (VAST 2010)*, Salt Lake City, Utah: 59–66
- Andrienko G, Andrienko N, Bremm S, Schreck T, von Landesberger T, Bak P, and Keim D 2010b Space-in-time and time-in-space self-organizing maps for exploring spatiotemporal patterns. *Computer Graphics Forum* 29: 913–22
- Crandall D J, Backstrom L, Huttenlocher D, and Kleinberg J 2009 Mapping the world's photos. In *Proceedings of the Eighteenth International Conference on World Wide Web*, Madrid, Spain: 761–70
- Goodchild M F 2007 Citizens as voluntary sensors: Spatial data infrastructure in the world of Web 2.0. *International Journal of Spatial Data Infrastructures Research* 2: 24–32
- Keim D A 2005 Scaling Visual Analytics to Very Large Data Sets. WWW document, <http://infovis.uni-konstanz.de/events/VisAnalyticsWs05/pdf/03DanielKeim.pdf>
- Kisilevich S, Keim D, and Rokach L 2010 A novel approach to mining travel sequences using collections of geotagged photos. In Painho M, Santos M Y, and Pundt H (eds) *Proceedings of the Thirteenth International Conference on Geographic Information Science*. Berlin, Springer-Verlag: 163–82
- Peuquet D J 1994 It's about time: A conceptual framework for the representation of temporal dynamics in geographic information systems. *Annals of the Association of American Geographers* 84: 441–61
- Shneiderman B 1996 The eyes have it: A task by data type taxonomy for information visualizations. In *Proceedings of the IEEE Symposium on Visual Languages*, Washington, DC: 336–43
- Slocum T A, McMaster R B, Kessler F C, and Howard H H 2009 *Thematic Cartography and Geovisualization*, 3rd edn. Upper Saddle River, NJ, Pearson Prentice Hall
- Tobler W 1987 Experiments in migration mapping by computer. *The American Cartographer* 14: 155–63