

This is a pre-copy-editing, author-produced PDF of an article accepted for publication in 'Journal of Location Based Services' following peer review.

Leveraging Explicitly Disclosed Location Information to Understand Tourist Dynamics: A Case Study

Fabien Girardin^a, Filippo Dal Fiore^b, Carlo Ratti^b and Josep Blat^a

^aDepartment of Information and Communication Technologies, Universitat Pompeu Fabra, Barcelona, Spain

^bDepartment of Urban Studies and Planning, Massachusetts Institute of Technology, Cambridge, USA

Abstract. In recent years, the large deployment of mobile devices has led to a massive increase in the volume of records of where people have been and when they were there. The analysis of these spatio-temporal data can supply high-level human behavior information valuable to urban planners, local authorities, and designer of location-based services. In this paper, we describe our approach to collect and analyze the history of physical presence of tourists from the digital footprints they publicly disclose on the web. Our work takes place in the Province of Florence in Italy, where the insights on the visitors' flows and on the nationalities of the tourists who do not sleep in town has been limited to information from survey-based hotel and museums frequentation. In fact, most local authorities in the world must face this dearth of data on tourist dynamics. In this case study, we used a corpus of geographically referenced photos taken in the province by 4280 photographers over a period of 2 years. Based on the disclosure of the location of the photos, we design geovisualizations to reveal the tourist concentration and spatio-temporal flows. Our initial results provide insights on the density of tourists, the points of interests they visit as well as the most common trajectories they follow.

Keywords: spatio-temporal data analysis; geovisualization; location-disclosure; location-based services

1. Introduction

In the past years, research in location sensing and tracking has been dominated by figuring out where persons and objects are in space (see Hazas et al., 2004 for a review). The wide adoption of the mobile and wireless technologies often referred to as Location-Based Services allows people a new perception of their surrounding physical and social spaces. In parallel, the records of where and when people have been (i.e. spatio-temporal data), produced by these services, have led to improve the understanding of different aspects of mobility and travel. The analysis of these data helps to recognize the modes of mobility (Sohn et al., 2006), define significant places (Ashbrook and Starner, 2002), cluster tourist routes in a city (Asakura and Iryob, 2007), infer travel purposes (Wolf, 2001) or predict a driver's destination as a trip progresses (Krumm and Horwitz, 2006). Ratti et al. (2006) benefit from people's experience of mobile devices to gain a more thorough understanding of urban environments, and this is the field of application undertaken in this paper.

First, let us remark that the large spatio-temporal logs can be seen as personal, as we explain later and their analysis can provide urban planners, local authorities and designers of location-based services with information on how a city gets used by different groups complementing current data sources, which are normally used as the basis of decision and policy making. Knowing who populates different parts of the city at different times can lead to the provision of customized services (or advertising), the rescheduling of monuments opening times or the reallocation of existing service infrastructures. From a quite different perspective, city users themselves could be aware of the current ways in which they populate the city, and adopt different strategies as a result.

Our approach takes advantage of the recent explosion in the use of capture devices (e.g. mobile phones, digital cameras) and collaborative web platforms to share their content (see Torniai et al., (2007) for a review). This people-generated information provides large amounts of digital data linked to the physical world. Recent research showed the potential of the geographically annotated material available on the Web. For instance, Ahern et al., (2007) and Snavely et al. (2006) developed means of world exploration via photos and maps to foster "virtual tourism". Using a similar dataset, Rattenbury et al. (2007) showed that the location and time metadata associated with photos and their tags enable the extraction of "place" and "event" semantics. In our study, we focus on the sense of presence extracted from a collection of photographs captured in trips. We consider that uploading, tagging and disclosing the location of a photo can be interpreted as an "I was here" statement indicating the physical presence in space and time. By extension, we propose that people-generated geographically referenced information provide new insights on how people travel and experience the city.

The validation of the relevance of this concept takes body in the Italian Province of Florence. The local authorities of that region aim at better understanding the tourist flows travelling across the cities boundaries. So far, they have been using classical survey-based hotel and museums frequentation data to know where tourists of different nationalities prefer to spend their time, hence money. However, they lack observations of

This is a pre-copy-editing, author-produced PDF of an article accepted for publication in 'Journal of Location Based Services' following peer review.

the mobility, nationality and quantity of the “day trippers”, that is the tourists who visit Florence but are "invisible" in the data, as they do not sleep in town. In consequence, we retrieved 81017 photos taken in the region by 4280 photographers over a period of 2 years from the popular photo-sharing web platform Flickr¹. Based on the time and the disclosed location of the photos, we extracted records of the people presence and movements; performed statistical analysis and designed geovisualizations. This exploratory visual analysis was used as a mean of preliminary investigation. In fact, the results go far beyond the initial expectations of collecting clues on “day trippers” activities, as we shall see, providing new insights, and even a novel paradigm of urban geography.

In the remainder of this paper, we discuss first current work on mobility data collection and constraints to perform travel surveys. These shortcomings suggest that the research community should investigate and evaluate new data and perspectives. Then, we propose a novel approach that takes advantage of spatio-temporal data generated by tourists when publicly sharing their photos on the world-wide web. Next, we describe the types of data that can be collected and their meaning. Afterwards, we present the preliminary results of our analysis supported by geospatial visualizations. They highlight the ability to quantify the concentration of tourists and their movements over time through the major areas of interests. Finally, we conclude with a description of the implications from this case study and discuss the meaning for future work.

2. Related work and their limitations

The recent emergence of location technologies and techniques favoured the development of new approaches to capture and analyze people’s mobility (see Wolf, 2004 for a survey). The aim has been to replace traditional travel diaries, paper-and-pencil interview, computer-assisted telephone interviews, and computer-assisted-self-interview, by automatically collecting mobility data. For instance, Wolf et al. (2001) proved the feasibility of using Global Positioning System (GPS) data loggers to improve the quality or completely replace traditional questionnaires. However, this type of mobility survey faces the problem inherent to longitudinal studies such as recruiting a pool of respondents or preventing any fatigue effects. Besides the privacy concerns of continuously and precisely tracking people, Schoenfelder et al. (2002) and Stopher et al. (2003) identify the potential technical drawbacks of a GPS-based approach. They list transmission problems, warm-up times before getting a valid position, and the cost of post-processing the GPS data as issues that impair the quality of the survey. Indeed, Wolf et al. mention that the equipment packages deployed for their pilot study had a lot more problems than anticipated: the units and cabling used were not optimized for durability, resulting in the loss participants due to problems with equipment performance and user operations.

Other mobility studies relied on the mobile phones use of the Global System for Mobile communications (GSM) network to generate mobility data. In a first type of approach, the mobile devices calculate and report their position to a centralized service. The TeleTravel System (TTS) project (Wermuth, 2003) combined a mobile device GSM

¹ <http://www.flickr.com>

This is a pre-copy-editing, author-produced PDF of an article accepted for publication in 'Journal of Location Based Services' following peer review.

tracking technology and an electronic travel diary to determine the travel behaviour of the respondents; while Asakura and Iryob (2005) determined it by interpreting changes in the set of nearby towers and signal strengths of the phones as indicative of position and motion of their owners; Froehlich et al. (2006) augmented the GSM mobility data with context-triggered in-situ survey in which panelists rate the place they are in to study travel routines. A second approach employs the measurement by mobile network operators of traffic intensity and migration of each cell in a GSM network to capture the movement patterns of mobile phone users. Through this technique Ratti et al. (2006) measure the evolution through space and time of the activities in a city to support urban planning. Based on comparable network data, Ahas et al. (2007) retrospectively link the digital track of visitors with visited events and locations retrospectively. While this approach scales because it does not rely on a costly software or hardware deployment, it fails in capturing individual traces. A third approach relies on the deployment of Bluetooth enabled systems. Recent studies have been able to establish the flow of people at strategic locations (O'Neill et al., 2006) as well as to recognize daily user activities and to identify socially significant locations (Eagle and Pentland, 2006).

However, there are several key issues when using technologies to collect travel behaviours. A major concern consists is the privacy and ethical issues related to collecting data without individual's consent (Gutman and Stern, 2007). Indeed, there is an increase in the risk of identifying people or organizations when their data spatial precision improves. Other issues include the length of the survey (e.g. to prevent fatigue effects); the ability to collect individual traces; and the scalability challenge of deploying the tracking system, for instance, when dealing with a variety of cellular network standards and providers. Each of the above-mentioned approaches has to face these aspects. Table 1 shows how GPS, GSM and Bluetooth tracking perform with respect to these important aspects of a travel survey.

Table 1. Mobility data capture techniques with respect to issues in the context of travel surveys. Low=little issue, High=big issue

| Mobility data capture techniques / Issues | Scalability | Longevity | Individual traces | Privacy |
|---|-------------|-----------|-------------------|---------|
| GPS | High | High | Low | Medium |
| GSM (device-based) | High | Medium | Medium | Medium |
| GSM (aggregated network-based) | High | Low | High | Medium |
| Bluetooth | High | Low | Low | Medium |

Some scholars work in the combination of the approaches to produce more complete spatial data (origin, route, destination) when tracking individual's routes (Kracht, 2004). However, while promising, the approach has not been able to address all the issues. In the remainder of this paper, we argue that a new type of explicitly disclosed spatio-temporal data coming from public web platforms can overcome these constraints and provide an additional insights to understand the dynamic of people in an urban space.

3. The raw data and the processing

We study photos uploaded on the popular photo-sharing platform Flickr. People use this service to share and organize photos; an option allows to add them geographical references. Each time a photo is virtually linked to a physical location, the Flickr system assigns a longitude and latitude and retrieves the time of capture from the Exchangeable Image File Format (EXIF) metadata embedded in the photo. The location provided by the user generally indicates where the photo was taken; but sometimes it denotes the photographed object. When the user provides the location of the photo, the zoom level of the map (from 8 for region/city level to the most precise 16 for street level) is recorded as an accuracy attribute, completing the spatio-temporal information.

Figure 1 describes the process of recording and collecting the data. First, tourists take photos during their trips and journeys. Later, they manually associate a position to the photos through a Flickr map interface or other external map-based services. A minority of tech savvy users has their photos implicitly annotated with a position extracted from data collected by GPS devices embedded or external to the camera. From the publicly available photos in a given area, we retrieve the coordinates, timestamp, accuracy level, and an obfuscated identifier of the owner via the Flickr API. Based on these data, as an example, a chronologically ordered set of photos gives rise to traces that reveal the movements of different individuals in space. In the figure we show as well some analysis tools we discuss later.

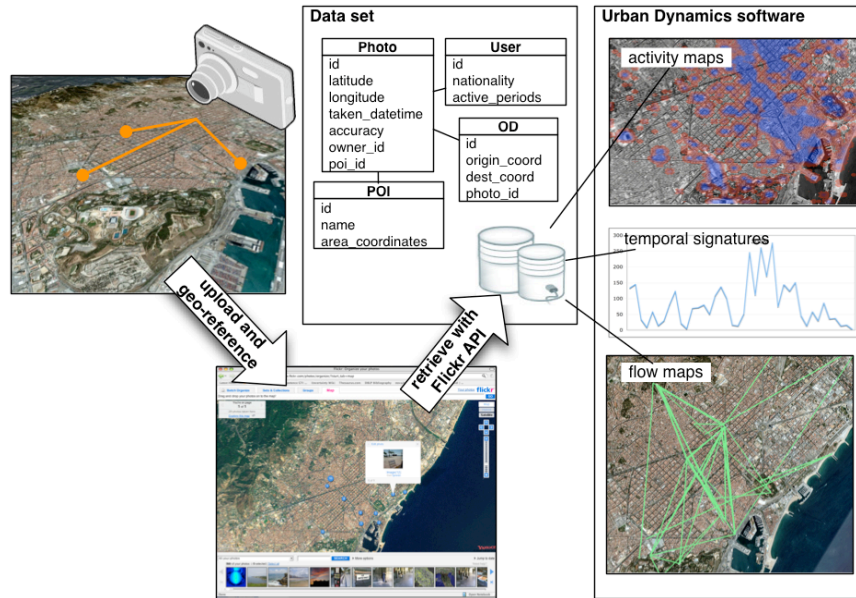


Figure 1. Data flow, from data recording, retrieving, storing to the visualizations. Photographic imagery copyright 2007 DigitalGlobe.

Our data set contains photos from the Province of Florence in Italy over a 2-year period, from November 2005 to November 2007. The timestamps extracted from the camera-generated EXIF metadata do not necessarily match the correct time. Indeed, the capturing devices might not be set with the correct time zone information, and the accuracy of the time embedded in the photo when it is taken is uncertain. However, the temporal analysis we describe later indicates that this seems to have little impact on our analyses. For example, as we consider only photos stamped in 2005, 2006 or 2007, we can assume that the user had set the camera date as its default value is rarely one of these years.

A first step of pre-processing was to separate visitors from residents of the region. To achieve that, we used the presence in the area over time as the discriminating factor. We divided the time in 30-day periods and computed the number of periods each photographer was active in the area. If a photographer took all his/her photos within 30 days, the algorithm considers him/her as a visitor, while if there is an interval greater than 30 days between two photos taken, he/she is categorized as resident. The aim of this strict threshold was to capture the real one-time tourists. Out of a population of 4280 photographers, we identified 3505 one-time visitors.

In addition, we were interested to know more about the nationalities of our photographers. Hence, we took advantage of the Flickr social function that invites (without forcing) users to provide “offline bits” on their city and country of residence. We found out that 65% (2782 out of 4280) of the users actually disclosed this information. While it is hard to predict how much of this data is truthful, the country of residence could be retrieved by automatically parsing the data in most cases. We had to assign manually a country in 11% (306/2782) of the cases, because of spelling errors or idiosyncratic names (e.g. “Big Apple” for New York).

This is a pre-copy-editing, author-produced PDF of an article accepted for publication in 'Journal of Location Based Services' following peer review.

Before going any further, let us remark that collecting and analyzing any kind of mobility data raises serious privacy issues; people are concerned about revealing the history of their whereabouts to un-trusted third party applications. We are concerned as well and our approach addresses this at two different levels. First, the users explicitly disclose the position of their photos on maps and control who gets access to their location data for our data set. Moreover, while obtaining this public information, we applied an obfuscation algorithm to lose the relationship with the identity used in Flickr. Thus, we only analyze anonymized records of digital traces publicly disclosed by individuals.

Our experience of collecting Flickr geo-referenced data provided some insights on the data at hand, which are an indicator of different aspects of their value. As of March 2007, Flickr contained more than 20 million photos linked to a physical location. Cities such as London or New York contain more than 250'000 photos and 9'000 single photographers each and growing pace of around 400 photos a day for London, and 150 a day for Barcelona. This quantitative richness might push towards an even bigger increase of publicly accessible people-generated location and time of the Flickr data set. At another level, a very specific type of people use Flickr to geographically reference their images. They are generally well travelled and technologically savvy. Therefore, we are dealing with a very specific type of tourists.

In addition, unlike the automatic capture of traces, the manual location disclosure embedded in the act of geotagging of photo provides additional qualities. Positioning photo on a map is not simply adding information on its location, it is also an act of communication containing what people estimate as being relevant for themselves and others. In that sense, a specific richness of this dataset arises in the intentional weight people put in disclosing their photos. We show that they have a tendency to select the highlights of their discovery of the city and discarding the downtimes.

4. Tools used, and initial results

This section of the paper presents the geovisualizations that are produced by our tools and discusses their meanings. In the research process, we were driven by an overarching research question: what is the spatio-temporal behaviour of tourists from different nationalities, inside the Province of Florence? Flickr seemed an adequate context to extract data from, as it contains the spatial and temporal elements, as well as the nationality. However, we have to take into account the following caveats: we have already discussed the lack of temporal accuracy; second we cannot assume perfect mirroring of the Flickr spatio-temporal patterns and the actual ones, as we do not know where users have actually been between a given photo and the following one. In other words, we can only deal with approximations. With these limitations in mind, we provide an analysis of the data based on map visualizations to reveal the tourist concentration and flow within the Central Italy region and the city of Florence. Each set of maps aims at giving indications on the applicability of the collected data set to better understand the tourist dynamics within a certain area.

Our initial results bring a new perspective on two aspects related to spatial and temporal density and movements of visitors of the Province of Florence and its capital the city of

This is a pre-copy-editing, author-produced PDF of an article accepted for publication in 'Journal of Location Based Services' following peer review.

Florence: (a) Characterizing the areas of the city/region where the tourists are concentrated, and (b) Revealing spatio-temporal signatures: activity by day of the week and month of the year, and days of the year. In addition, two types geovisualizations seek to understand better the flow of tourists into, out of and within the Province of Florence: (c) trajectories into and out of the area studied, (d) Patterns of flow within the Province of Florence.

4.1. The “Urban Dynamics” software

The extraction of structure, meaning and insight from large, multifaceted, spatio-temporal datasets is a challenging task that requires skills not possessed by many engaged in geovisualization (Dykes et al., 2007). Our approach takes advantage of open and freely-available resources and combines them using de facto standards often based on the extensible Markup Language (XML). We use Google Earth² for interactive visual synthesis of encodings generated using a combination of MySQL³ for data storage and querying to select and aggregate, and a software developed in Java we named “Urban Dynamics” to access, process, transform, aggregate, cluster, sample, filter the raw data stored in the database and to generate outputs. The Keyhole Markup Language⁴ (KML) is used to describe visual encodings and define interactions. Google Earth is used as a means of interactively visually analysing and inspecting data, through its spatial and temporal navigation tools, its access to wider contextual data, and its ability to create animations with image overlays. Figure 2 describes the data processing and multiple outputs generated by Urban Dynamics. The data of the photos are stored in a spatial matrix to perform population density analysis and in linked arrays of positions to reveal the traces left by people in their visits. In practical terms, a trace consists in a chronologically ordered set of geographically referenced photos taken by one person over one day. Is this the appropriate place to define the trace? Should it be defined earlier? These data structure can be divided in temporal periods to ground the base of a spatio-temporal analysis and visualization. The density values stored in the spatial matrix are employed to cluster the main areas of interest in the studied area. After their storage in the database, these areas are mapped in Google Earth with a KML file for identification and labelling. The software also generates from the matrix data heat maps of population density that can be included as an image overlay in KML files visualized in Google Earth. We describe next the spatio-temporal geovisualizations of flows and density generated with Urban Dynamics and the analyses of their results.

² <http://earth.google.com/>

³ <http://www.mysql.com/>

⁴ <http://earth.google.com/kml/>

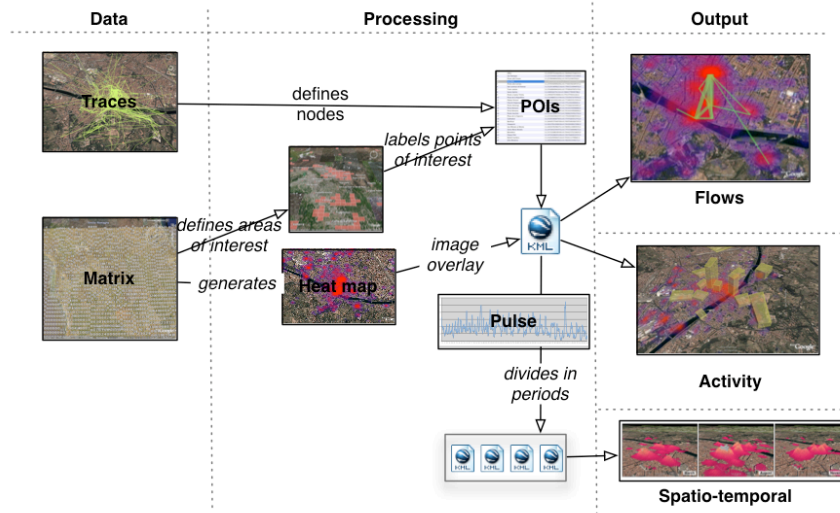


Figure 2. The Urban Dynamics software multiple data, processing and outputs. Photographic imagery copyright 2007 DigitalGlobe.

4.2. Where are the tourists concentrated?

To understand where tourists concentrate, we extracted from our dataset the latitude and longitude of the photos taken by visitors that were geotagged with a granularity more precise than the city resolution level. These data were put in a matrix covering the studied area with each cell of the matrix containing the number of photos and visitors ever present in that zone. To reveal visually the tourist concentrations, we proceeded in two steps. First, we generated from the matrix a Keyhole Markup Language (KML) document to feed it to the Geocommons⁵, a geospatial data visualization platform. It produced an interactive interface with heat maps of the tourist concentrations on top of the Google Maps system (Figure 3). The results show a zoomable map of the overall tourist activity covering the northern part of Central Italy, the city of Florence and around the famous Basilica di Santa Maria del Fiore.



Figure 3. Interactive zoomable map of the overall tourist activity from the northern part of central Italy to fine-grained information on the Basilica di Santa Maria dal Fiore in Florence. Generated with GeoCommons. Photographic imagery copyright 2007 Cnes/Spot, NASA, DigitalGlobe and TerraMetrics.

These visualizations show that the main activity in the Province of Florence actually takes place in the city of Florence. There are other major points of interest in the region, such as the cities of Pisa and Siena, as well as the Mediterranean coast between La Spezia

⁵ <http://www.geocommons.com/>

and Livorno, and the Island of Elba. Within the city, the tourists tend concentrate around the Ponte Vecchio and the Basilica di Santa Maria del Fiore. While these geovisualizations give a quick and simple overview of the presence of tourists, they do not offer a scale to understand the quantitative meaning of the colours. Therefore, as a second step, we developed an algorithm with the Processing⁶ Java Library to generate heat maps for their inclusion as an additional layer in Google Earth. Next, we clustered the presence values in the matrix to define the major areas of tourist concentration. A list of 50 areas each representing a point of interest was computed and manually labelled according to their position on the map. The match of the areas with the matrix values allowed rank the major points of interest of the area studied according to the number of photos taken, the number of tourists and residents, number of photos taken by person and the likeliness to encounter residents or tourists (Figure 4). The Basilica di Santa Maria del Fiore is the most visited monument of the city followed by the famous Ponte Vecchio, the Plaza de la Signoria and Palazzo Vecchio. However, the Arco Lorena is the monument that tourists like to photograph the most (18.06 photos per tourist) compared to the Basilica di Santa Maria del Fiore that triggers 7.72 photos per tourist. Fiesole, a town in the suburbs of the city is one of the main point of interest fairly out of the tourist frenzy with a proportion of visitors of 73% compared to the 89% of the Ponte Vecchio.



Figure 4. Presence of tourists in the main areas of interest in downtown Florence. The visualization shows areas of activity (i.e. photos taken) in red and the presence of tourists with yellow polygons. The altitude of the polygons represents the number of individuals present. Photographic imagery copyright 2007 DigitalGlobe.

4.3. Temporal activity and spatio-temporal signatures

In our first attempt, we produced maps of tourist concentrations without taking into account the period when data were recorded. In this section, we generated similar maps to compare the activity in time, namely, seasons and weekends versus weekdays, etc. As differential data might be more informative, we took the “tourist pulse” of the area of study, by developing first its spatio-temporal signature. We charted the temporal

⁶ <http://www.processing.org>

signature of the number of visitors active in Figure 5, which shows the weekly, monthly and yearly pulse in 2007 with the number. Unsurprisingly, the summer months are the busiest of the year and tourists are more active on Saturdays and Sundays. A yearly vision of the presence reveals a steady pulse with peaks on the weekends.




Figure 5. Pulse as the average number of visitors per each month of the year, day of the week and day of the year for 2007.

Based on the 2006 data, we produced animations of monthly activity to understand where and when was taking place the tourist activity (Figure 6). Our geovisualization adds a layer on top of a map of the Province, showing 250 x 250 cells of 20m x 12m whose colors represented the number of photos taken and visitors present. Similarly to Real-time Rome animations based on Erlang values (Reades et al., 2007), we used a 5 persons threshold to reveal the main areas of activity. That is, any cells of the matrix containing less than 5 persons were discarded. In addition we smoothed the visualization of areas of activity with an interpolation algorithm. This type of animation helped revealing the areas and monuments popular at certain seasons. The Boboli gardens, for instance, do not draw many visitors in winter while other points of interests do not lose much attractiveness.




Figure 6. Screenshots of the spatio-temporal animation of the presence of tourists in downtown Florence in 2007. Photographic imagery copyright 2007 DigitalGlobe.

4.4. Inbound and outbound trajectories

As we discussed earlier, the classic survey-based hotel and museums frequentation data used by tourist authorities to know where tourists of different nationalities come from and go do not capture “day trippers”, the tourists who visit Florence but do not sleep in town. Our data brought a new perspective on the issue to understand where the visitors had been prior to entering the area and where they were heading after leaving it. We retrieved the images the 3505 visitors took prior to entering, and after leaving the area. As a result, we collected more than 1.8 million geographically referenced photos taken worldwide. Then, we recovered the inbound and outbound trajectories through the rules described in

Table 2. A specific problem was to find the proper threshold in time to infer a precise origin of an inbound movement and the destination of an outbound one. The test of different thresholds of 24, 48 and 72 hours revealed that more than 83% of the inbound and outbound movements took place within 24 hours. The movements exceeding this interval were discarded, as they might not have reflected the proper origin and destination of a journey via Florence.

| Type of trajectories | Algorithm |
|----------------------|--|
| Inbound | Detect the position and time of the last photo before entering the area, and the position and time of the first photo after entering the area. |
| Outbound | Detect the position and time of the first photo after leaving the area, and the position and time of the last photo before leaving the area. |

With the similar process and tools as the ones described in section 4.1. we plotted two types of geospatial visualizations to provide information on where tourists are prior to entering the city of Florence and after leaving it. The first result is an interactive zoomable heat maps that reveals density of inbound and outbound destination of a day trip in Florence. With the manual annotation of the heat maps we were able to define the 28 main areas and cities of origin and destination. Each area had a bounding box and assigned with the number of incoming and outgoing visitors computed. With this quantitative understanding of the trajectories, our software generated flow maps to offer another perspective on the movement by rendering the trajectories of tourist. They are generated from KML files of origin and destination traces loaded in the Google Earth software. For flows leading out of Florence (Figure 7), each green trace links the position of the last photo taken in Florence with the position of the first photo taken outside the city within 24 hours. Figure 7b displays through an analogous strategy, the aggregates outbound movement from the city of Florence. The comparison of inbound and outbound trajectories shows striking similarities. The most connected cities, Rome, Pisa and Venice generated a relatively equal amount of incoming and outgoing movements. Florence acts as connecting city between the north and the more southern part of Italy with Rome as the main destination.

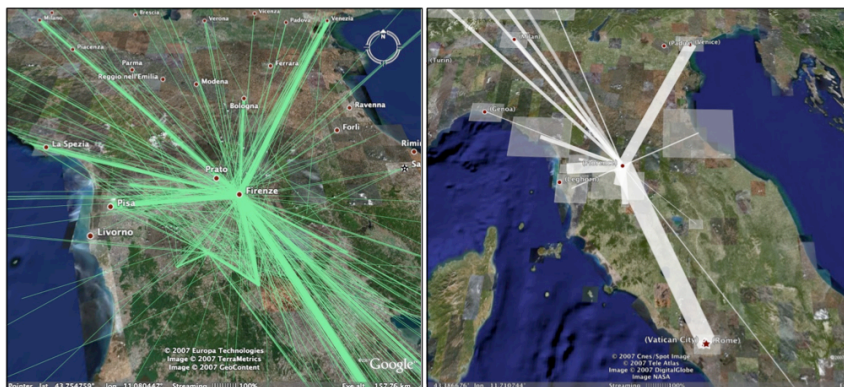


Figure 7. Outbound trajectories from the city of Florence. The map displays the individual traces in green

and their aggregation in white. Photographic imagery copyright 2007 Cnes/Spot, NASA, DigitalGlobe and TerraMetrics.

4.5. Patterns of flow

In addition to revealing the origins and destinations of the tourists, the analysis of the dataset allows to gather insights on the flows within the boundaries of the northern part of Central Italy and the city of Florence. We could trace the Flickr users from the digital footprints they leave along their path. Aggregating these personal footprints and formatting them in KML reveals the travel behaviours of specific types of visitors. For instance, Figure 8 shows that tourists from the USA follow a specific graph constituted by the nodes of Florence, Siena, Pisa, Genova and Perugia. On the other hand, Italians are more adventurous in their exploration of the area. For instance it shows the popularity of the Island of Elba, an attraction not visited by American tourists.

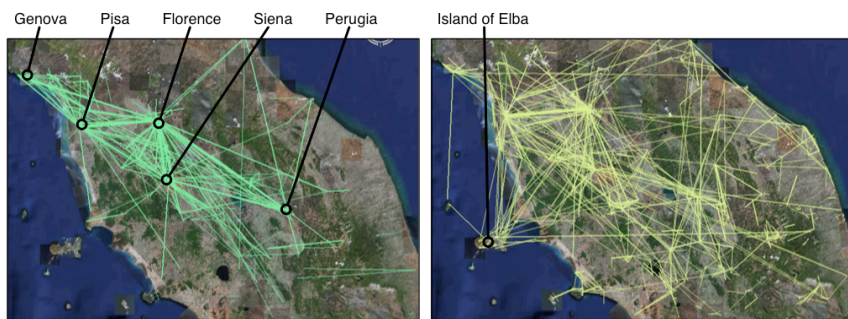


Figure 8: Movements of Americans in green compared to Italians in yellow in northern part of central Italy. The comparison of the visualization shows the distinct was to apprehend the space. Photographic imagery copyright 2007 Cnes/Spot, NASA, DigitalGlobe and TerraMetrics.

In practice, our basic flow maps can easily become too cluttered and prevent the detection of flow patterns. As a solution, we use the main areas of interests of the Province of Florence and the city. Our visualization software matches these areas of interests with the traces left by visitors to gain a quantitative measurement of the movements between the main attractions of the area studies. Figure 9 shows the result of this process for downtown Florence. Our software generates the traces left by each visitor and aggregates them in correspondence to the main points of interests they connect. A KML file is generated with the heat map as an image overlay object and the flows between each point of attraction plotted as green lines with a weight proportional to the density of the flow. This combined vision of the density and flow of tourists shows that they consume a very limited space of the city. The tour to the unique relatively distant area of interest, the viewpoint of the Piazzale Michelangelo, seems to take place from and the most popular monument of the city, the Basilica di Santa Maria del Fiore.

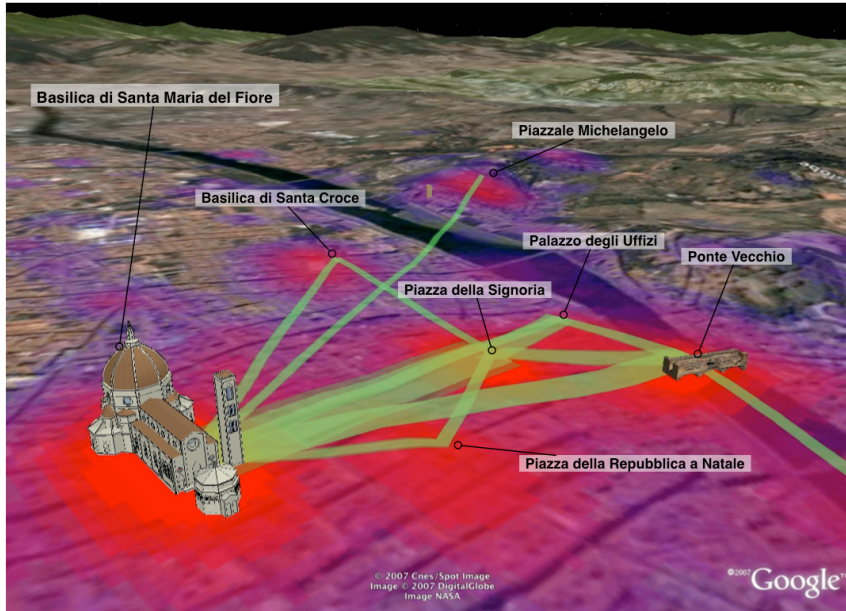


Figure 9: Downtown Florence, main areas of tourist activity (in red) and flows of visitors between them (in green). The width of the green line represents the amount of visitors who moved from one attraction to another. Photographic imagery copyright 2007 Cnes/Spot, NASA, and DigitalGlobe.

5. Conclusion and future work

The explosion in the use of captures devices (e.g. mobile phone, digital cameras) and the emergence of content sharing platforms is leading to the emergence of a wealth of publicly available user-generated geospatial data. Our case study specifically featured the value of Flickr and its geographically reference photos in the Italian Province of Florence with the goal of performing urban and mobility analysis. Our approach takes advantage of individual traces made publicly available to analyze behaviours in a urban space. Their exploratory analysis enables us to quantify by the amount of photos taken and the presence of individuals the attractiveness over time of the major points of interests of an urban space. In addition, the mapping of these data shows the flows entering, leaving and visiting a city like Florence. Not only the photos allow us to quantify the movements between attractions, they also help detecting the distinct patterns of mobility among groups of tourists of different nationalities. This type of insight is limited in most travel surveys and urban sensing infrastructure by privacy-sensitive or aggregated information.

The careful processing of the user-generated data, and the use of appropriate geovisualizations are necessary ingredients of the analyses performed. We have indicated where some improvements of the processing and the visualizations might come from. For instance, we need to find out more on the representativeness of the user of Flickr and their veracity in disclosing the photos they share. User-generated data require special attention to self-selection in the data sample. Indeed, those people who proactively upload digital information on open platforms such as Flickr are very likely to represent a tech-savvy sub-section of any given population. By collecting further data on the socio-demographics of these users, it might be possible to understand better how much the

This is a pre-copy-editing, author-produced PDF of an article accepted for publication in 'Journal of Location Based Services' following peer review.

behaviour findings from research based on a self-selected sample can be generalized to a given population.

This case study could open a new perspective in urban cartography and lead to a new urban paradigm based on the analysis of publicly available people generated geo-referenced data. The explicit act of sharing this information with the ability to control and cloak the data greatly reduces privacy concerns inherent to current travel surveys techniques, as we have shown in this paper. In addition, unlike the classic automatic capture of traces, the act of geotagging a photo can be interpreted as an act of communication because people only give information they estimate as being relevant for themselves and others. In that sense, a richness of people-generated geographically referenced data relies on the specific effort people put in to disclosing the location information. Our geovisualizations show that they have a tendency to select the highlights of their discovery of the city and discard the downsides.

In the future, similar types of spatially anchored user-generated content might surface to become relevant for travel, mobility and urban studies. Sources might range from implicit data from the usage of radio-frequency networks (GPS, GSM, WiFi, Bluetooth) to more explicit information contained in geospatial web applications (e.g. geo-referencing in Flickr), and the emerging social applications based on the disclosure of presence and location information. Researchers, urban planners, local authorities, and others might consider the aggregate analysis of these different channels.

As part of further work, we consider several extensions to the promising results presented in this paper. First we aim at correlating our dataset with other spatio-temporal data such as GSM network usage, hotel and attractions surveys. Second, we believe that the recording of such data could be used to inform the design and deployment of location-based services to enhance the tourist experience. So far location-aware applications have tended to concentrate on using a mobile carrier's immediate geographic location. We believe that using a position history to tailor results from requests for information should enhance them (Mountain and Raper, 2001). This work shows that there might be potential in taking advantage of the digital traces people constantly leave behind them to, for instance, reveal the temporal character of a space, its attractiveness among a certain group of people, and its level of connectivity with other spaces.

Acknowledgements

The results reported in this paper are part of a broader research undertaken in the Wireless City, Florence and Beyond project. We would like to thank the authorities of the Province of Florence for their support. Also, we are indebted to many people at the Massachusetts Institute of Technology and the Universitat Pompeu Fabra for providing extremely stimulating research environments and for their generous feedback. In particular, thanks to Assaf Biderman, Francisco Calabrese, Nicolas Nova, Bernd Resch and Toni Navarrete for letting us pick their brains. Of course, any shortcomings are our sole responsibility.

This is a pre-copy-editing, author-produced PDF of an article accepted for publication in 'Journal of Location Based Services' following peer review.

References

- Ahas, R., Aasa, A., Silm, S., Tiru, M. 2007. Mobile positioning data in tourism studies and monitoring: case study in Tartu, Estonia. In: Sigala, M., Mich, L., Murphy, J. (Eds.), *Springer Computer Science: Information and Communication Technologies in Tourism*, ISBN: 978-3-211-69564-7, pp. 119-128.
- Ahern, S., Naaman, M., Nair, R., and Yang, J. H. 2007. World explorer: visualizing aggregate data from unstructured text in geo-referenced collections. In *JCDL '07: Proceedings of the 2007 conference on Digital libraries*, pages 1–10, New York, NY, USA. ACM Press.
- Asakura, Y. and Iryob, T. 2007. Analysis of tourist behaviour based on the tracking data collected using a mobile communication instrument. *Transportation Research Part A: Policy and Practice*, 41(7):684–690.
- Ashbrook, D. and Starner, T. 2002. Learning significant locations and predicting user movement with gps. In *ISWC '02: Proceedings of the 6th IEEE International Symposium on Wearable Computers*, page 101, Washington, DC, USA. IEEE Computer Society.
- Dykes, J., Slingsby, A., and Clarke, K. 2007. Interactive visual exploration of a large spatio-temporal dataset: Reflections on a geovisualization mashup. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1176–1183.
- Eagle, N. and Pentland, A. S. 2006. Reality mining: sensing complex social systems. *Personal Ubiquitous Comput.*, 10(4):255–268.
- Froehlich, J., Chen, M. Y., Smith, I. E., and Potter, F. (2006). Voting with your feet: An investigative study of the relationship between place visit behavior and preference. In *Ubicomp*, pages 333–350.
- Gutman, M. and Stern, P., editors 2007. Putting People on the Map: Protecting Confidentiality with Linked Social-Spatial Data. *National Academies Press*.
- Hazas, M., Scott, J., and Krumm, J. 2004. Location-aware computing comes of age. *IEEE Computer*, 37(2):95–97.
- Kracht, M. 2004. Tracking and interviewing individuals with gps and gsm technology on mobile electronic devices. In *Seventh International Conference on Travel Survey Methods*.
- Krumm, J. and Horvitz, E. 2006. Predestination: Inferring destinations from partial trajectories. In *Ubicomp*, pages 243–260.
- Mountain, D. and Raper, J. 2001. Modelling human spatio-temporal behaviour: a challenge for location-based services. In *6th Internat. Conference on GeoComputation*. University of Queensland, Brisbane, Australia.
- O'Neill, E., Kostakos, V., Kindberg, T., gen. Schieck, A. F., Penn, A., Fraser, D. S., and Jones, T. 2006. Instrumenting the city: Developing methods for observing and understanding the digital cityscape. In *Ubicomp*, pages 315–332.
- Ratti, C., Pulselli, R. M., Williams, S., and Frenchman, D. 2006. Mobile landscapes: Using location data from cell-phones for urban analysis. *Environment and Planning B: Planning and Design*, 33(5):727 – 748.
- Reades, J., Calabrese, F., Sevtsuk, A., and Ratti, C. 2007. Cellular census: Explorations in urban data collection. *IEEE Pervasive Computing*, 6(3):30–38.

This is a pre-copy-editing, author-produced PDF of an article accepted for publication in 'Journal of Location Based Services' following peer review.

- Rattenbury, T., Good, N., and Naaman, M. 2007. Towards automatic extraction of event and place semantics from flickr tags. In *SIGIR '07: Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 103–110, New York, NY, USA. ACM Press.
- Schoenfelder, S., Axhausen, K., Antille, N., and Bierlaire, M. 2002. Exploring the potentials of automatically collected gps data for travel behaviour analysis - a swedish data source. In Möltgen, J. and Wytzisk, A., editors, *GI-Technologien für Verkehr und Logistik*, number 13 in IfGIprints, pages 155–179. Institut für Geoinformatik, Universität Münster.
- Snaveley, N., Seitz, S. M., and Szeliski, R. 2006. Photo tourism: exploring photo collections in 3d. *ACM Trans. Graph.*, 25(3):835–846.
- Sohn, T., Varshavsky, A., LaMarca, A., Chen, M. Y., Choudhury, T., Smith, I., Consolvo, S., Hightower, J., Griswold, W. G., and de Lara, E. 2006. Mobility detection using everyday gsm traces. In *Ubicomp*, pages 212–224.
- Stopher, P., Bullock, P., and Jiang, Q. 2003. Visualising trips and travel characteristics from gps data. *Road & Transport Research*, 12(2):3–14.
- Torniai, C., Battle, S., and Cayzer, S. 2007. Sharing, discovering and browsing geotagged pictures on the web. Technical report, HP Labs.
- Wermuth, M., Sommer, C., and Kreitz, M. 2003. *Impact of New Technologies in Travel Surveys*, chapter 27, pages 455–482. Pergamon, Oxford.
- Wolf, J. 2004. Applications of new technologies in travel surveys. In *7th International Conference on Travel Survey Methods, Costa Rica*.
- Wolf, J., Guensler, R., and Bachman, W. 2001. Elimination of the travel diary: experiment to derive trip purpose from global positioning system travel data. *Transportation Research Record*, (1768):125–134.