

# Photographer Paths: Sequence Alignment of Geotagged Photos for Exploration-based Route Planning

**Abdallah El Ali**  
ISLA, University of  
Amsterdam  
Amsterdam, The Netherlands  
a.elali@uva.nl

**Sicco N. A. van Sas**  
University of Amsterdam  
Amsterdam, The Netherlands  
sicco@ddo.nl

**Frank Nack**  
ISLA, University of  
Amsterdam  
Amsterdam, The Netherlands  
nack@uva.nl

## ABSTRACT

Urban mobility analysis of geotagged photos can unlock mobility patterns of users who took these photos, which can be used for exploration-based city route planners. Applying sequence alignment techniques on 5 years of geotagged Flickr photos in Amsterdam (The Netherlands) allowed creating walkable city routes based on paths traversed by multiple photographers (or photographer paths). To evaluate our approach, we conducted a user study with Amsterdam residents to compare our routes with the most efficient and popular route variations. Drawing on experience questionnaire data, web survey responses, and user interviews, our results show that our photographer paths were perceived as most stimulating and suitable for city exploration. Moreover, while digital aids based on photographer paths can potentially aid city exploration, their acceptance in mainstream route planners likely depends on their visualization. From our proof-of-concept approach and user study findings, we discuss the potential of data-driven exploration-based city route planners.

## Author Keywords

Exploration-based route planning; sequence alignment; geotagged photos; UGC; urban computing

## ACM Classification Keywords

H.5.2 Information interfaces and presentation: Theory and methods; H.4 Information Systems Applications

## INTRODUCTION

We are not always in a hurry to get from point A to point B. Sometimes we take a longer route because it is more scenic, more interesting, or simply to avoid the mundane [12]. While expert tour guides (e.g., Lonely Planet<sup>1</sup>) tell us what to see and do, they are geared towards recommending destinations and tour guide offers, not generating route plans or journeys. In fact, research has focused extensively on tourism, and replete with how to develop mobile technology (or electronic mobile guides) to support travelers and tourists in ‘what to

do’ [13, 3]. For current route planning services (e.g., Google Maps<sup>2</sup>), the generated routes are tailored towards providing shortest paths between any two locations. However, city pedestrians, whether tourists or locals, may not always want the fastest route – this is strengthened when for example considering the buzz surrounding Foursquare’s<sup>3</sup> launch of the Explore functionality that recommends places based on friends’ check-ins. In other words, given the right context and time, people do wish to wander into hitherto unfamiliar or unconventional paths. However, there is surprisingly little work in CSCW and geographic HCI [11] that addresses this gap: how can we build city route planners that automatically compute route plans based not on efficiency, but on people’s trailing city experiences? Importantly, how do these experiences influence our route preferences and perception of urban spaces?

With the unbridled adoption of location-aware mobile devices that permit geotagging multimedia content, places and routes can now be ubiquitously micro-profiled with geotagged user-generated content. This geotagged data comes from mobile social media services (e.g., Flickr<sup>4</sup>, Twitter<sup>5</sup>), and relates to the actions and experiences of thousands of people at different locations. In line with a recent SIG meeting discussing the research opportunities of geographic HCI and the rise and use of User-Generated Content (UGC) [11], we believe this geotagged data can be used not only for revealing the social dynamics and urban flow of cities [15], but also unlock fragments of user intentions and experiences at places and transitions between them. This data can provide a latent source for generating exploration-based routes traversed in a city that are not based on travel efficiency.

In this paper, we focus on sequences of geotagged photos, which we show can allow computing city paths that represent the history of where the photographers of these photos have been. By using this latent information on photographer paths, we believe this unlocks novel application and research avenues for data-driven exploration-based route planners.

## RESEARCH QUESTIONS

Our main research question is: how can we *automatically* generate walkable route plans in a city for residents that would like to *explore a city*? Specifically, what existing data sources and which methods can be used to generate such

<sup>1</sup><http://www.lonelyplanet.com/> ; last retrieved: 08-11-2012

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CSCW’13, February 23–27, 2013, San Antonio, Texas, USA.

Copyright 2012 ACM 978-1-4503-1209-7/13/02...\$15.00.

<sup>2</sup><http://www.maps.google.com/> ; last retrieved: 08-11-2012

<sup>3</sup><https://foursquare.com/> ; last retrieved: 08-11-2012

<sup>4</sup><http://www.flickr.com/> ; last retrieved: 08-11-2012

<sup>5</sup><http://www.twitter.com/> ; last retrieved: 08-11-2012

routes, and how are these routes perceived in comparison with both the popular and fast routes in a city? While our target user group is city residents (defined as having lived in the city for at least one year), our contributions as will be shown later also apply to tourists who wish to discover off-beat paths when visiting a city.

To define what constitutes an interesting walkable route, we reasoned that the mobility behavior of city photographers tells us something about worthy alternative routes in a city. The underlying assumption here is that locations of photographs are potentially interesting as the photographer(s) found it worthwhile to take a picture there. For this purpose, the image photo-sharing site Flickr provides a suitable data source given that many images are geotagged.<sup>6</sup> Here, we focus on users that do not have any specific interest or do not want to supply this interest and they just want to be given an interesting route from A to B, which we suspect city photographers (be they locals or tourists) can help unravel.

To avoid making the user supply preferences, we wanted to automatically generate routes based on where people traveled within a city. But not every route may be interesting, so we focused on routes made up of locations where people took pictures, given our assumption that taking a picture somewhere depicts an interesting location. One such route made out of photographs from a single photographer is insufficient, so ideally we want multiple photographers that took pictures at the same locations in the same sequence, i.e. took the same route and found similar things photo-worthy along the same locations. Thus, we needed a method that handles not only where photographers have been, but importantly, in what order they have been there and to what extent their movements resemble the movements of other city photographers. To achieve this, we use sequence alignment (SA) methods. These methods are borrowed from bioinformatics and later adapted to time geography to systematically analyze and explore the sequential dimension of human spatial and temporal activity [20]. We hypothesize that the aligned routes traversed by multiple city photographers (or ‘photographer paths’) provide desirable paths for pedestrians wishing to explore a (familiar) city. Furthermore, while we are concerned with route planning using both mobile devices and desktops, here we focus on pre-trip route plans, which usually involves viewing routes on a desktop.

Our work yields two main research contributions: a) a novel data-driven methodology for generating walkable route plans based on photographers’ paths in a city and b) an empirical understanding (based on quantitative and qualitative assessments) of how users perceive these photographer paths in comparison with today’s efficiency driven route planners and popular routes. Additionally, we provide a preliminary investigation on the role that digital information aids on a map (e.g., Points-of-Interest (POIs), photos, comments, etc.) play in influencing people’s decisions about which route to take for exploring a city.

<sup>6</sup>Around 520,00 geotagged photos tagged with ‘Interesting’ in Amsterdam alone (retrieved on 30-05-2012).

The rest of the paper is structured as follows: we give a review of related work, followed by our Photographer Paths approach and alignment experiments. We then present a user study (consisting of a lab and web-based study) to evaluate the different route plans and importance of digital information aids in influencing users’ perception of city routes. We then present and discuss our results, future work and conclude.

## RELATED WORK

Given our interest in both generating and consuming UGC-generated routes, this paper draws from various related work, including time geography, urban modeling techniques, and importantly route planners.

### Time Geography

Time geography dates back to Hägerstrand [9], who stressed the importance of taking into account temporal factors in spatial human activities. This gave rise to a space-time path visualization which shows the movement of an individual graphically in the spatial-temporal environment when one collapses the 3D space and uses perpendicular direction on a 2D map to represent time.

Essentially, time geography seeks to analyze patterns of human activity using space-time paths in an objective, structural manner (e.g., aligning sequences of activities by visitors to the Old City of Akko [20]). The idea behind this is to visualize human movement and interactions between individuals on a 2-D plane where the x- and y- axis represent geographic coordinates (longitude and latitude, respectively) and the z-axis represents time. This so-called space-time “aquarium” is used for analysis and evaluation of social dynamics and activity distribution across space and time. This is useful for analysis of aligned sequences of human activity, where the activity of concern here is the photo-taking behavior by photographers of the geotagged images retrieved from the Flickr photo-sharing site. In short, we use these representational methods to analyze sequences of photo-taking activities, where we later use alignments for generating walkable city routes based on these photographer paths.

### Photo-based City Modeling

Given the iconic correspondence between photographs and reality, we believe photo sharing services like Flickr provide a window into the unique perspectives of city photographers. If we consider Flickr photo features, thousands of photos shared by photographers come contextualized with textual user-defined tags and descriptions, geotags (latitudes and longitude coordinates), and time-stamps (date and time of day).

[21] used the varied photos taken by multiple photographers of the same scene along a path as controls for image-based rendering, allowing automatic computation of orbits, panoramas, canonical views, and optimal paths between 3D scene views. Relatedly, [23] used a game-based crowdsourcing approach to constructing 3D building models, based on contributions from a community of photographers around the world. In this work however, we are not concerned with 3D

scene views (e.g., Google Street View<sup>7</sup>), only with the generation and perception of route plans plotted on a 2D map.

Using Flickr data alone, computational approaches have been developed to understand tourist site attractiveness based on geotagged photos [8], constructing travel itineraries [7] and landmark-based travel route recommendations [16], and generating personalized Point-of-Interest (POI) recommendations based on the user's travel history in other cities [6]. All these approaches however focus primarily on describing locations and/or landmarks at these locations, and not on within-city routes that connect them irrespective of popular landmarks. Closer to the present approach, [18] mine sequences of locations from Flickr geotags – however, their focus is on recommending the most popular tourist places in a city.

### Non-efficiency Driven Route Planners

Relevant here is whether there is work on route planners that go beyond finding routes that optimize commute efficiency. [17] developed a system to automatically generate travel plans based on millions of geotagged photos and travel-ogues, which was tailored towards providing city tourists with popular attractions/landmarks and popular routes between them. Relatedly, [4] mined people's attributes from photos to provide personalized travel route recommendations; however, their method was aimed at finding personalized hotspots, not for exploring off-beat paths in a city.

[1] categorized travel trips from people based on geotagged images taken and the accompanying tags and photo titles, allowing development of an application for searching frequent trip patterns. While the goal here was catering for users that wish to learn more about the most frequently visited places in a city, we are interested in automatically computing route plans for exploring a city based on sequences of photographers' movements. Relatedly, [12] present a method comprising a user survey and subsequent clustering analysis to classify route selection criteria for bicyclists. Here, they found that bicyclists most favored fast and safe routes, followed by simple and attractive ones in an urban environment. Finally, using a crowdsourcing approach, [26] developed the Mobi system which allows people to collaboratively create and edit itinerary plans in cities, thus showing the merits of human computation tasks to provide rich plans. In our work however, we try to automate the process of providing exploration-based route plans in a city.

## PHOTOGRAPHER PATHS

### Approach: MSA of Arbitrarily Long Sequences

To align the geotagged photos, we used the ClustalTXY [25] alignment software. ClustalTXY is suitable for social science research, as it makes full use of multiple pairwise alignments, where alignments are computed for similarity in parallel – in other words, it makes use of a progressive heuristic to apply multiple sequence alignment (MSA) [25]. Furthermore, ClustalTXY allows representing up to 12-character words, which allows us to uniquely represent small map regions containing the geotagged photos.

MSA is done in 3 stages: first, pairwise alignments are computed for all sequences. Then these aligned sequences are grouped together in a dendrogram based on similarity. Finally, the dendrogram is used as a guide for multiple alignment. To deal with differences in sequence length, ClustalTXY adds gap openings and extensions to sequences. Opening is the process of adding a gap between two previously gapless words and extension is the process of adding another gap in between two words which already had a gap.

Throughout the paper, 'words' are synonymous with 'locations' and 'nodes', where a given term is used depending on the context of discussion. The more aligned sequences that contain the same words, the more popular is a particular word. Thus, the most interesting sequences are distilled by finding matching sequences of popular words in the alignment results. In our approach, we map each location in a sequence to a cell in a partitioned grid map where each cell corresponds to an indexed location unit (e.g., 125 x 125 m cell). For example, a route containing 5 locations would thus be BcEeFsgQlQn, where Bc constitutes the first word (i.e., a location). Furthermore, all repeated words were trimmed down to one (e.g., FyEjEjEjYfWyFs would become FyEjYfWyFs). We use a simple grid-based approach instead of a mean-shift clustering approach (cf., [5]) in order to allow for locations photographers visited that may not otherwise contain many data points. We then apply MSA to the photographer routes (consisting of sequences of their photos' locations) to find the aligned location sequences. These are used for selecting matching segments of sequences across photographers – we call these exact matches *photographer route segments* (PRSs).

### Dataset

We used the Flickr API to retrieve geotagged photos within Amsterdam, The Netherlands (17.3 km N-S and 24.7 km E-W)<sup>8</sup> over a 5-year period (Jan. 2006 - Dec. 2010), with the following attributes: owner ID, photo ID, date and timestamp, tags, latitude, longitude and the accuracy of the coordinate. This resulted in a database of 426,372 photos.

### Preprocessing

We included in our database only photos with geo-coordinates with near-street accuracy or better (accuracy 14-16 in Flickr attributes). We inferred the sequences taken by photographers from the time and geotags of their photos. Each photo in the sequence had to be taken within 4 hours from the previous photo. Sequences were constrained to having at least two or more different locations (or nodes), where each location corresponds to a cell on the grid. Given early experiments, we used a grid cell size of 125 x 125 m. We also now focused our grid on the city center of Amsterdam as most routes were in this area and this would speed up alignment computation. The city center could be described using a grid of 26 by 26 cells, so 2-letter words were sufficient. These steps resulted in a dataset of 1691 routes, which had an average length of 9.92 words (min = 2, max = 124). There were

<sup>8</sup>The area is based on the Amsterdam region as currently defined in the Flickr API (bounding box: 4.7572, 52.3178, 5.0320, 52.4281; centroid: 4.8932, 52.3731).

<sup>7</sup>[www.google.com/streetview/](http://www.google.com/streetview/) ; last retrieved: 08-11-2012

1130 unique photographers, where on average each photographer contributed 1.50 routes to the dataset.

**Sequence Alignment**

Main parameters in MSA are gap opening and extension values. In bioinformatics these values correspond to a penalty for splitting a DNA or protein sequence, which needs to be restricted in order to retain informative groups of sequences of nucleotides or amino acids. In our case this analogy does not hold and we want to match as many words as possible, therefore we set both values to 0. Alignment for this 125m dataset took approximately 7 hours on a single core server.

To find photographers paths from PRSs, we set constraints for selecting PRSs having *at least 4* photographers having *at least 2* aligned nodes (or locations/words). This choice was motivated by the resulting PRSs from our 5-year dataset (see Fig. 1), where having more photographers per route segment took precedence over number of locations (or nodes). These 2 or more aligned nodes form the PRSs. Photographers could have made different photos in between nodes, but they must have visited the locations in the same order and within 4 hours between each visited location. After applying these constraints, we had 231 PRSs (visualized in Fig. 2) with an average length of 2.61 nodes and a maximum PRS length of 4 nodes.

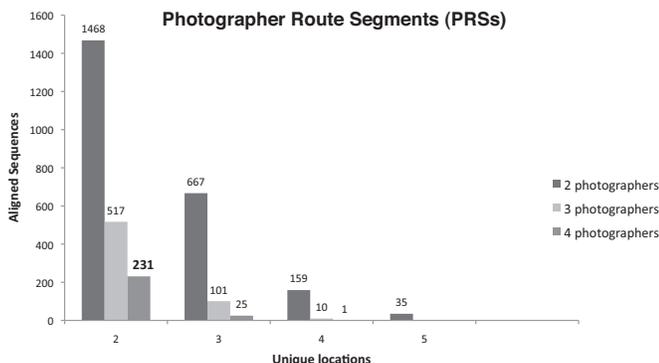


Figure 1. Aligned sequences (PRSs) in Amsterdam over a 5-year period for different numbers of unique photographers and locations. Our PRS set choice value ('231 sequences') is shown in bold.

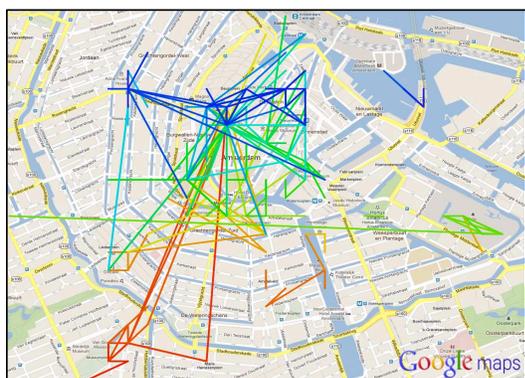


Figure 2. 231 PRSs of alignments of 4 photographers and 2 unique locations in Amsterdam city center. Best seen in color.

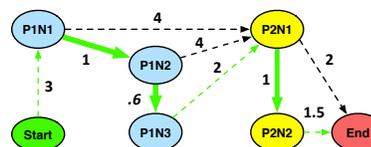


Figure 3. Example of PRS aggregation. 'P1N2' stands for node 2 of PRS 1. Numbers indicate inter-node distance. Best seen in color.

**PRS Aggregation**

Next step was to develop a method which uses these PRSs to generate routes from a given start location to a user specified destination. We used an implementation of Dijkstra's shortest path algorithm<sup>9</sup> to find the shortest route along the PRSs. Recall that a PRS is a transition between two or more locations/nodes based on sequences of at least 4 photographers, where a PRS is calculated from the aligned sequences of the ClustalTXY alignment.

Dijkstra's algorithm requires a network of edges between nodes, with a specified cost for traversing each edge. We thus had to specify how our PRSs would both connect within themselves and to each other. Every edge cost is set to the distance between the nodes. However, if all nodes were to be connected with each other, then Dijkstra's algorithm would simply output the direct connection between the start node and the end node as a route, so we set a maximum distance for edges between and within PRSs. Dijkstra's algorithm finds the shortest path between nodes, but we wanted to steer the algorithm to make use of as many transitions between nodes within each PRS as possible, even if this meant a detour, because these transitions are more representative of the actual paths of photographers. To solve this, we required that at least two nodes were used in each PRS, thus at least one edge within a PRS is always used. After this hard constraint, Dijkstra's algorithm connects a PRS to another PRS, because using a third node within the original PRS will usually result in a longer route. The final route would thus be made out of PRSs where only two nodes within each PRS are used. To maximize the number of nodes within each PRS, we gave discounts [range 0-1] to the distances of every edge (beyond the first edge) used within a PRS, forcing Dijkstra's algorithm to incorporate extra edges within the PRS.

A simplified PRS aggregation task using these methods is shown in Fig. 3. The thick solid lines show the edges between the nodes within PRSs, while the thin dashed lines show the connections between the PRSs and the user specified start and end nodes. Dijkstra's algorithm would normally find the following shortest path Start-P1N1-P2N1-End with a cost of 9, but due to the constraint of at least two nodes per PRS and the discounted edge cost (0.6 (cost) \* 1 (original weight) = 0.6; shown in italics) between P1N2-P1N3, a different route is selected. The recommended route (green edges or Start-P1N1-P1N2-P1N3-P2N1-P2N2-End) will now make use of all the PRS edges.

We applied this algorithm on our chosen PRS set (4 photographers, 2 locations), to create two different photographer

<sup>9</sup><http://code.activestate.com/recipes/119466-dijkstras-algorithm-for-shortest-paths/>; last retrieved: 08-11-2012

routes in the city center of Amsterdam: one made up of 9 PRSs with 11 total connections (where black route segments are gaps filled for completing the route), that runs from Central Station to Museumplein (CM). The other was made up of 4 PRSs with 6 total connections (again black route segments are route gaps filled), and runs from Waterlooplein to Westerkerk (WW). These routes are visualized in Fig. 4.

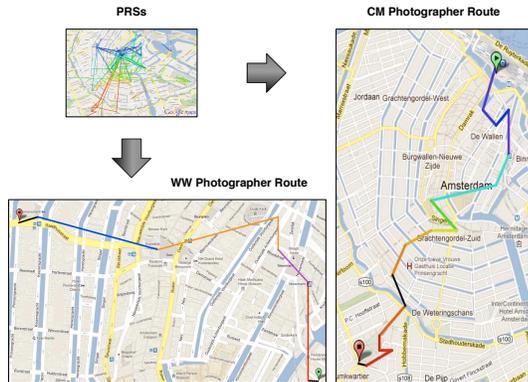


Figure 4. Our chosen PRS set after applying the modified Dijkstra’s algorithm resulted in two ‘crude’ photographer routes (where individual PRSs are color coded): a) Central Station to Museumplein (CM) route b) Waterlooplein to Westerkerk (WW) route. *Best seen in color.*

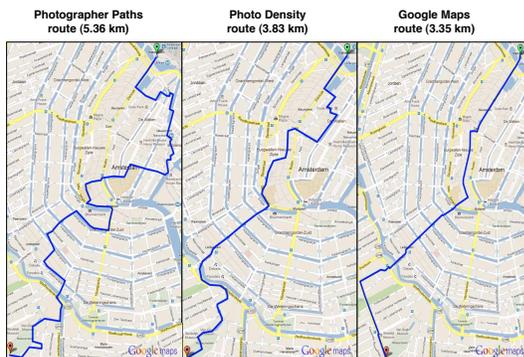


Figure 5. Visual comparison of the generated routes from Central Station to Museumplein. *Best seen in color.*

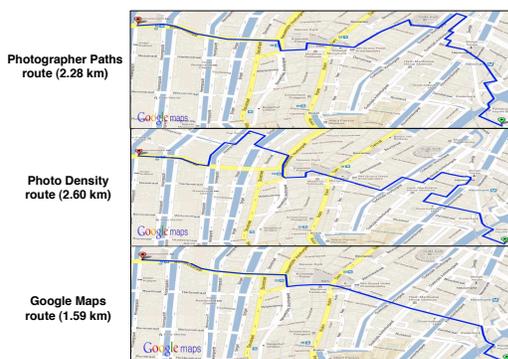


Figure 6. Visual comparison of the generated routes from Waterlooplein to Westerkerk. *Best seen in color.*

**Results**

To turn the ‘crude’ photographer routes given by our adapted Dijkstra’s algorithm into walkable routes, these routes were mapped to a Google Maps map where we took the shortest walking distance between each route node (or location). This resulted in walkable photographer paths. The Photographer Paths (PP), Photograph Density (PD), and Google Maps (GM) route variations for our chosen two routes, Central Station to Museumplein (CM) and Waterlooplein to Westerkerk (WW) route are shown in Fig. 5 and Fig. 6, respectively. Explanation and motivation for the PD and GM route variations is given below.

**USER EVALUATION**

**Laboratory-based study**

*Study Design*

We wanted to evaluate whether our Photographer Paths (PP) route variations are indeed preferred by users for exploration-based route planning. While previous work (e.g., [14]) addressed how to evaluate the usability of electronic mobile guides (which may include route planners), there are no established standards on how to best evaluate a service that provides alternative walkable city routes from a human-centered perspective, especially since POI selection accuracy and routing efficiency are not suitable measures for the desirability of the service. However there is work that addresses similar problems. [19] evaluated location-based stories generated automatically from Wikipedia<sup>10</sup> by making use of a Likert-type questionnaire. [22] evaluated their mobile guide in a cultural heritage setting by means of participant observation, questionnaires, and semi-structured interviews. [16] used a quantitative approach where they compared their photographer behavior model against three probabilistic models to account for the accuracy of their personalized route recommendations.

To evaluate whether our route variations are perceived by users as desirable alternatives to current efficiency-based route planners, we chose a user-centered mixed-methods approach that includes both quantitative and qualitative measures. While our target user group included both city residents and tourists, here we focus on expert evaluations from city residents. To test whether our approach provides not just a novel method for generating routes, but routes that city residents would rate as preferable for an exploration scenario, we chose to compare our generated route variation with two other route variations that have a similar start and end destination. Our baseline comparison was a route based on the density of photographs taken per grid cell, where we assumed that this would provide a route plan through the most touristic hotspots within Amsterdam. This was chosen instead of a route that connects a density of all POIs as a POI-density based route would require further differentiating between the kinds of POIs, which is not the aim of a route planner that generates routes automatically without requesting user preferences. For each scenario, participants (who were city residents) had to evaluate the routes in Amsterdam.

Two routes were tested, each with 3 variations: Photographer Paths (PP) route, a Photo Density (PD) variation as baseline,

<sup>10</sup><http://www.wikipedia.org/> ; last retrieved: 08-11-2012

and a Google Maps (GM) efficiency-based route variation. For each route, participants were given scenarios. For the first scenario, participants had to imagine being in the company of local friends on a sunny Saturday between 14-15:00 o'clock, where they wished to walk from Central Station to Museumplein (CM route). For the second scenario, participants had to imagine being in the company of a friend (a local) on a cloudy, Sunday evening between 19-20:00, where this friend just returned from a vacation and they now wished to catch up at a café near Westerkerk (WW route). While both scenarios emphasized there was time to spare, we expected participants to favor efficiency in the WW route.

PD route was created by drawing a path between grid cells containing the highest density of geotagged photos taken in Amsterdam in 5 years, for the hour corresponding to each scenario given to participants (14-15:00 and 19-20:00, respectively). The restriction by hour was set so paths between cells remains meaningful, as plotting a 5-year dataset of geopoints makes it difficult to differentiate between choosing one cell over another. This route served as a popular and touristic route baseline by which to measure our PP route against.

We set up a within-subjects experimental design, where route variation is the independent variable, and measured dependent variables are: a) perceived quality of the presented route variations for each route (CM and WW) b) participants' route preferences and c) subjective reports on what they thought about the generated routes. To measure perceived quality of the route variations, we adapted the AttrakDiff2<sup>TM</sup> [10] questionnaire<sup>11</sup> so that participants can reflect on the presented routes and give us a quantitative measure of the hedonic and pragmatic aspects of each route variation. AttrakDiff2 measures pragmatic and hedonic qualities by allowing participants to provide ratings on a 7-point semantic differential scale for 28 attributes<sup>12</sup>, resulting in 4 quality dimensions: 1) Pragmatic Quality (PQ), which measures usability of a product (or in our case routes). Here, PQ gives insight into how well users can achieve their goal given each route 2) Hedonic Quality - Identification (HQ-I), which gives insight into the extent that users can identify with a given route 3) Hedonic Quality - Stimulation (HQ-S), which gives insight into the extent that a route stimulates users with novelty and enables personal growth 4) Attractiveness (ATT), which provides a global value and quality perception of a route, or in other words, perceived attractiveness.

To get further insight into participants' perception of the route variations, we had a two part semi-structured interview at the end of each testing session, where users could give their feedback directly on their route preferences and inform us what information types they find valuable in deciding whether or not a route affords exploration. For the first part, participants were asked to give their opinion on which route variation they

preferred for each route, and what they thought about routes based on photographer paths. In the second part of the interview, they were provided with examples of different digital information aids and asked which (or a combination of) they found useful for exploring a city. These were: a plain Google Maps route, Foursquare POIs (that include short textual comments left by others) along a route, a route with Flickr photos, our PP route segments (made up of PRSs) that shows via color coding the different route segments that make up the photographer paths (see Fig. 4), and a route showing multiple photo geopoints (i.e., PD route). Finally, they were asked about the applied potential of this kind of route planning service.

The need to investigate information types (even if not the primary aim of our study) relates to the need for transparency and intelligibility in ubiquitous computing systems [24]. On one hand, to make a fair user perception comparison between routes generated by route planners means that further information cannot be provided on a route variation. This is because we risk comparing the effects of information type on route preference, and not the quality of the route itself. On the other hand, in an actual route planner system, users should be given the option to understand 'why' a given route is generated, which is why we had to simultaneously investigate digital information aids in different kinds of media.

#### Participants

15 participants (10 male, 5 female) aged between 21-35 ( $M_{age} = 29.2$ ;  $SD_{age} = 3.3$ ) were recruited. Only participants who had lived in Amsterdam for at least one year were recruited, to ensure that they were able to adequately judge the presented route variations. Our participant sample spanned 9 different nationalities. Most participants claimed to know Amsterdam fairly well (10/15), where the rest knew it either very well (2/15) or just average (3/15). Many (10/15) had a technical background (e.g., Computer Science), and all were familiar with route planning services, where most (10/15) reported using route planners at least once a week.

#### Prototype, Setup & Procedure

To test the route variations, an interactive web-based prototype route planner interface was shown to each participant. The interface was adaptable to mobile devices, but testing route preferences on a mobile device was not important as participants were selecting a route based on pre-trip preferences, which usually involves viewing routes on a desktop. The study was conducted at the User Experience lab at XYZ university. Each session took approximately 45 min. to complete. To facilitate discussion and eventual consensus regarding our interview questions amongst participants, participants were interviewed in groups of three. For the first part of the study, each participant was seated in front of a laptop, where they each interacted (zooming, panning) with the route planner interface. For the interview, participants were allowed and encouraged to discuss and answer the questions in a collaborative manner.

Before the study session, each participant filled a background information form, signed an informed consent form, and read through detailed instructions for performing the task. In each task, a participant had to inspect 3 route variations (PP, PD,

<sup>11</sup>AttrakDiff2<sup>TM</sup> is a questionnaire originally developed to measure the perceived attractiveness of interactive products based on hedonic and pragmatic qualities. However, the measured bipolar qualities that apply to interactive products can also apply to city routes, making for a suitable domain generalization.

<sup>12</sup>Only one attribute-pair was changed to fit our study: Technical-Human was replaced with Slower-Faster for the PQ dimension.

GM) for each route (CM, WW). The order of presentation of both routes and route variations were counterbalanced and then randomized (after first presented variation) across participants. After inspecting each route variation, a participant had to fill in the ‘same’ adapted AttrakDiff2™ questionnaire, marking their responses with the corresponding route variation number so their responses were relative to one another. Participants were asked to give their first, spontaneous response. After all three participants finished inspecting all route variations, they were given the semi-structured interview. After the interview, each participant was offered a small monetary reward and thanked for their time.

**Web Survey**

To test if there is an immediate difference between the generated route variations (PP, PD, and GM), we constructed a short web-based survey to compare each of the route variations for the CM and WW routes. This survey was meant to be short and easy to fill, and to provide a rough idea of whether people consider route plans based on alternative city routes useful, and to collect data on what they find useful digital information aids for exploration. Basic demographic information (age, gender, years in Amsterdam, familiarity with route planners) was asked, and thereafter participants had to choose which route variations they would follow given our two respective scenarios (as in the lab-based study). Here however, there were two main differences to the lab study: a) maps showing each route were static images, and so participants could not zoom in on a location and b) all routes and route variations were shown on a single screen, so order effects were not accounted for.

For the survey, 82 participants (55 male, 27 female) aged between 17-62 ( $M_{age} = 27.6$ ;  $SD_{age} = 6.1$ ) responded. Most (44/82) lived in Amsterdam for more than 3 years, some (15/82) between 1-3 years, and the rest either less than a year (11/82) or only visited Amsterdam before (12/82). All were familiar with route planners and GPS-based systems. All participants were considered here, as we were interested in immediate reactions to the presented variations and questions about digital information aids.

**Results**

*Perceived Route Quality*

Responses on the modified AttrakDiff2 in the lab study were analyzed within groups, per generated route. For each category, one-way repeated measures ANOVA tests were conducted comparing results from all route variations. Means, standard deviations, confidence intervals, significance, and (partial) eta-squared values for each tested route variation for each route (CM and WW) are shown in Table 1. For the CM route, a repeated measures ANOVA showed significant differences in the responses across quality dimensions for only PQ, HQ-S, and ATT. For the WW route, a repeated measures ANOVA showed significant differences in the responses across quality dimensions for only PQ and HQ-S. Post-hoc pairwise comparisons (with Bonferroni correction<sup>13</sup>) between each variation (PP, PD, GM) were conducted in every case.

<sup>13</sup>Backward-corrected SPSS© Bonferroni adjusted p-values are reported.

Dimension	Route Variation	M	SD	CI	P-value	$\eta_p^2$	
CM Route	<b>PQ</b>	<b>PP</b>	-1.5	.1	[-2,-1]	<b>p&lt;.001</b> $F(2,28) = 38.9$	.7
		<b>PD</b>	.4	1.0	[-1,-.9]		
		<b>GM</b>	1.7	1.1	[1.2,2.2]		
	HQ-I	PP	.5	.8	[.1,.9]	p = .262	.1
		PD	.4	.7	[.04,.7]	$F(2,28) = 1.4$	
		GM	-.01	.9	[-.5,.4]		
	<b>HQ-S</b>	<b>PP</b>	1.2	.8	[.8,1.6]	<b>p&lt;.001</b> $F(1.3,19.4) = 21.3$ (corr. G-G $\epsilon = .69$ )	.6
		<b>PD</b>	.2	.9	[.1,.3]		
		<b>GM</b>	-1.1	1.1	[-1.7,-.6]		
	ATT	<b>PP</b>	.9	1.2	[.4,1.6]	<b>p&lt;.05</b> $F(2,28) = 4.8$	.2
		<b>PD</b>	.8	1.0	[.3,1.3]		
		<b>GM</b>	-.4	1.4	[-1.3]		
WW Route	<b>PQ</b>	<b>PP</b>	-.9	1.4	[-1.6,-.2]	<b>p&lt;.001</b> $F(2,28) = 38.18$	.7
		<b>PD</b>	-1	1.2	[-1.6,-.4]		
		<b>GM</b>	2.3	.9	[1.9,2.8]		
	HQ-I	PP	.1	.6	[-.2,.4]	p = .591	.04
		PD	.3	.5	[.05,.5]	$F(2,28) = .5$	
		GM	.3	.7	[-.01,.7]		
	<b>HQ-S</b>	<b>PP</b>	.5	1.0	[-.01,1]	<b>p&lt;.001</b> $F(2,28) = 17.3$	.5
		<b>PD</b>	.7	.9	[.2,1.3]		
		<b>GM</b>	-1.4	1.1	[-2,-.9]		
	ATT	PP	.3	1.1	[-.3,.8]	p = .877	.01
		PD	.3	1.0	[-.2,.9]	$F(2,28) = 1.3$	
		GM	.5	1.0	[-.04,1]		

**Table 1. Descriptive statistics (N=15) for each route variation (PP, PD, GM) under each tested route (CM and WW) for each AttrakDiff2 dimension: PQ = Pragmatic Quality, HQ-I: Hedonic Quality-Identity, HQ-S: Hedonic Quality-Stimulation, ATT: Attractiveness.**

Where significant, dimensions are represented in bold, and where a particular pairwise comparison is not significant, the dimensions are in (additional) italics.

For the CM route, participants perceived clear differences across all route variations for the PQ. It is not surprising that the GM variation scored the highest here, given that the route is based on efficiency. Nor is it surprising that our PP variation scored the lowest, given the length of the route. However, it is surprising to see that the PD variation significantly differed from both the PP and GM variation, as it is close in length to the GM variation, yet still considered not pragmatic. This can be partially explained by the fact that the PD route variation runs through all the heavily touristic areas, which may not be very practical to take. For the PP and PD variations, responses for HQ-I were around zero, however there were no significant differences between any of the variations. This suggests that all participants identified with the route variations similarly, but perhaps overall it may be that those variations were not ones that our sample of participants identified with.

For the HQ-S, there were significant differences between all tested route variations, where our PP variation scored the highest. This finding is in line with our hypothesis that photographer paths can provide a stimulating and novel route variation in a city. The mean responses for the PD variation are around zero, which shows that for our participants who are living in Amsterdam, the popular aspect of the PD variation was not stimulating (but nor was it found to be as completely dull and boring as the GM route). Finally, there were significant differences in ATT between PP and GM and between PD and GM route variations, however not between PP and PD. The fastest GM variation was rated the lowest. Here, we would have expected our PP variation to be rated higher

than the PD variation, however it could be that participants found the PD variation attractive depending on the company they are with (given the scenario); as they later mention in the interviews, they could be with touristic friends in which case they would find the popular route variations more attractive for the sake of tourism.

For the WW route, participants perceived clear differences for the PQ dimension between the GM route variation and both our PP variation and the PD variation. The GM route variation was not surprisingly rated as the most pragmatic, however here the PD variation was on par with our PP variation. For the HQ-I, there were again no significant differences between any of the route variations, were all responses had a mean around zero. This again suggests that identifying with the generated routes was not important for our participants. For the HQ-S, the GM variation was again rated the lowest, however this time there were no significant differences between our PP variation and the PD variation. This was likely due to the overall short distance between Waterlooplein and Westerkerk, where little room was left for identifying stimulating qualities of the routes. Importantly, in line with the scenario, participants here did not value stimulating qualities of the route variations, as one going to a café with a friend back from vacation on a cloudy day is not a situation that affords a stimulating route variation. This is further confirmed by considering the ATT dimension, where the GM route variation was now rated higher than it was for the CM variation, which indicates that participants valued the short distance for the GM variation. However, given that there were again no significant differences, it seems to still be that all route variations are attractive enough to take, even under efficiency-driven scenarios.

#### Route Preference

All but one participant in the lab study found the scenarios quite realistic. The one participant objected that it would have been clearer if the gender of friends in the scenarios was given. In the lab-based study, after inspecting all route variations for each route, participants were asked which variation they would follow (if any) and why (convenient, interesting, scenic, or other response). For the CM route, most participants (9/15) chose to follow our PP route variation, where the rest picked the PD (4/15) and GM (2/15) variations. For all cases where participants picked the PP variation, they stated it was more scenic. This is in line with our hypothesis that the PP variation provides a favorable route to take in a city one has time to explore. During the interviews, participants mentioned that they found the PP variation attractive and suitable for the scenario (P3: “*I liked the second route variation [PP] given it was nice weather so perfect for exploration.*”; P12: “*One of the routes [PP] was long and took many detours, and I thought that was a very attractive route!*”).

Also some reported that the PD route variation (as well as the GM variant) to be very touristic, which they did not like (P5: “*I would obviously not go through the Kalverstraat as it is very touristic.. the first [GM] and last [PD] route took me through there, and I think that is not really Amsterdam.*”). Still, others found the PP route too long (P10: “*I thought*

*the [PP] route was too long, my feet would fall off!*”) and preferred something in the middle between the PP and GM variations (P14: “*If I want to see the city, I would go for the [PD] one.*”). In the web survey, most participants said they would follow the GM route (40/82), followed by the PD route (23/82) and the PP route (10/82), where the rest chose neither (9/82). The difference in findings here can be explained either because participants did not inspect the routes carefully or take the scenarios seriously, or perhaps more likely those that lived in Amsterdam already knew their own specific routes, which they take by commuting by bicycle and largely avoiding touristic areas (where both PP and PD route variations pass through). Unlike the lab study, the experimenter could not steer participants to stick with the scenario – this was evident by most comments left in the free-form box form (e.g., “*I would not easily walk these routes... who in amsterdam walks? ;)*”; “*I would use cycling routes.*”). Nevertheless, they highlight that alternative route plans may take time to be accepted as city exploration aids (alongside routes suggested by expert travel guides) in both familiar and unfamiliar cities.

For the WW route, most participants (10/15) in the lab study favored the GM route variation, followed by the PD variation (4/15) and in one case neither one. For all participants that picked GM, they stated it was more convenient. This was not surprising under the WW scenario, where it was a cloudy evening and meeting a friend returning from a vacation. Indeed, later in the interviews this was mentioned explicitly (P3: “*You are going for coffee so you just want to get there, unlike in the first [CM] scenario where it is a nice day and you have time.*”). Here, it was also mentioned that our PP route variation and PD both went off to the Red Light District, which may not be desirable for them (P11: “*Second [GM] route was very simple and straightforward... Red light part is more interesting for tourists.*”). In the web survey, most participants said they would follow the GM route (67/82), followed by the PD route (6/82) and the PP route (3/82), where the rest chose neither (6/82). Favoring the GM variation for the WW scenario is in line with the findings of the lab-based study.

#### Digital Information Aids for City Exploration

During the post-test interview, participants were asked about digital information aids they would find useful when deciding to take a route in a city for exploration purposes. They were asked explicitly about what they thought about information that tells you how many persons (particularly city photographers) took a given route segment over a certain time period (e.g., 1 year). Around half (8/15) found such information useful for exploring a city one already knows, some where not quite sure (4/15), some thought it depended on who those photographers were (2/15), and one thought this is not something for him. In the second part of the interview (after participants saw examples of different digital information aids), many (10/15) found our PP information type attractive (e.g., P4: “*Nice to find corners that I don't really know.*”). This was visualized by coloring the different segments of a route on a map and stating how many photographers took a segment over time (for us, 4 photographers through 2 locations over a 5 year period). However, many of those participants

(6/10) also stated that they would combine this information with photos (3/6) and POIs (3/6). While some mentioned that the PP information is useful for exploring a familiar city (P8: “[PP] information is useful if you are familiar with the city.”), still others thought it redundant (P3: “Hard to see why I’d use a map in a city I already know.”).

In the web survey, we gained further insight into what digital aids people find useful in helping them explore a city. These are (ordered by count of mentions): POIs along a route (51x), route distance (51x), comments along a route (ranked by highest ratings or recency) (24x), expert travel guides (22x), photos of route segments (17x), no digital aids (13x), number of photographers that took a given path over a time period (9x), number of photos along a route over a time period (9x). While we expected that established aids such as POIs and distance are useful indicators for planning personal routes, it was surprising that there were only few reports of the usefulness of photographer paths to guide exploring a city. This was also evident with photo counts along a route, which is also a novel information aid. This may be because participants are not familiar with such novel aids, especially since in the web survey (to avoid biasing route preference and save time), they were not provided with visualizations of how this information may be visualized. Considering the findings of the lab-study and web survey, it seems that photographer paths as city exploration aids have potential, but stating it without visualizing how it could intelligibly augment an alternative route plan (as was done in the lab-based study) may have influenced its immediate adoption by participants in the web study.

#### *Use of Non-Efficiency Driven Route Planners*

Finally, participants were asked what they thought about future route planning applications that generate alternative walkable routes in a city, both locally and abroad, and whether they would use them. Most participants (13/15) were positive about such applications (P7: “Yes. Helps you to explore more and discover different things, gives you another option.”; P4: “Yes, especially for a specific corner of a city.”), where the remaining two brought to question why they would use such apps in a city they already knew (P12: “I would like it more for a city that I don’t already know.”; P14: “Google Maps is enough.”). Responses on what these route planners should be based included: personalized route plans (e.g., in accordance with route travel history) and preset route plans (e.g., museum route, market route, etc.). Together, the responses from participants highlight that a new generation of route planners aimed for exploring cities is desirable.

## DISCUSSION

### Study Limitations

In this paper, we have presented a novel approach based on sequence alignment methods to generate exploration-based route plans in a small-sized city like Amsterdam. Our lab-based user test provided strong evidence that our generated route plans are of potential interest to city dwellers especially by way of hedonic stimulation, under an exploration-based scenario. However, while most participants mentioned they would follow our PP route variation in the CM scenario, this was not the case for the web survey. This brought to question

not only whether such a quick survey was able to sufficiently provoke reflection on the given scenarios and routes, but also that alternative city route plans may not be immediately accepted by both locals and visitors to a city. Furthermore, without providing intelligible explanations (cf., [24]) as was done in the lab-study (visualizing the digital aids on the generated route plans) for why a route was given may have made it difficult for potential users to make an informed judgment.

Another limitation is the real-world evaluation of the generated routes. While we have tried to tease out the differences between each route on participants’ route preferences in a lab-based and web-based study, we may not immediately generalize to how participants would react to such routes if used in real-time in an actual situation with a working prototype. A related issue is that the number of route variations presented is limited to two, where participants’ reactions may differ in such cases. Nevertheless, the findings from the lab-study provide strong evidence that photographer paths have potential for generating desirable route plans in the city, and importantly highlight the merits of an automatic data-driven approach based on geotagged photographs in a city.

### Towards Automatic Exploration-based Route Planners

We started from the question of how human sequential movement patterns can be leveraged as an indicator of interest, and attempted to answer this by relying on a data-driven approach that borrows SA methods from bioinformatics and time geography. Our approach goes against the established human-centered literature on catering for user needs, where typical mobile recommender studies (e.g., [2]) begin with distilling requirements from observing and interacting with users. Here, it can be argued that such a quantitative approach may drastically oversimplify our human needs for exploring a city. However, we wanted to compute routes based on paths taken by photographers in a city *automatically*, without burdening the user to tell us her desires.

Under the foregoing motivation, we have created an opportunity to consider a currently unused information type obtained from geotagged images to guide exploration-based route planning in a city: the number of city photographers that took a given route segment over a certain time period. With this proof-of-concept approach, we have shown it is possible to leverage social geotagged data to cater for the hard problem of automatically generating exploration-based route plans.

### FUTURE WORK & CONCLUSIONS

Future work should take the work presented here into the real-world, where a functional prototype that generates photographer paths can be deployed in situ on a large-scale. This would allow (crowdsourced) data collection on more route variations than the ones considered here, and allow for grid cell sizes that are determined by the size of the considered street network (and not fixed to a size like 125 x 125m as considered here). The resulting bigger sample size would better support the alignment-based method presented, and the merits of data-driven exploration-based route planners in general.

In this paper, we have presented a proof-of-concept approach that uses sequence alignment methods and the digital foot-

prints of city photographers (obtained through Flickr) to compute exploration-based route plans within a small-sized city like Amsterdam. From our user study with Amsterdam residents, we found that our photographer paths are promising for city exploration. We hope to have set the stage for further experimentation with data-driven approaches to tackle the hard problem of automatically generating ‘interesting’ route plans for exploring both familiar and unfamiliar cities.

#### ACKNOWLEDGEMENTS

Authors thank Eurostar (E5262) SmartInside Project for its support.

#### REFERENCES

- Arase, Y., Xie, X., Hara, T., and Nishio, S. Mining people’s trips from large scale geo-tagged photos. In *Proc. MM ’10*, ACM (2010), 133–142.
- Bellotti, V., Begole, B., Chi, E. H., Ducheneaut, N., Fang, J., Isaacs, E., King, T., Newman, M. W., Partridge, K., Price, B., Rasmussen, P., Roberts, M., Schiano, D. J., and Walendowski, A. Activity-based serendipitous recommendations with the magitti mobile leisure guide. In *Proc. CHI ’08*, ACM (2008), 1157–1166.
- Brown, B., and Chalmers, M. Tourism and mobile technology. In *Proc. ECSCW ’03*, Kluwer Acad. (2003), 335–355.
- Cheng, A.-J., Chen, Y.-Y., Huang, Y.-T., Hsu, W. H., and Liao, H.-Y. M. Personalized travel recommendation by mining people attributes from community-contributed photos. In *Proc. MM ’11*, ACM (2011), 83–92.
- Clements, M., Serdyukov, P., de Vries, A., and Reinders, M. Finding wormholes with flickr geotags. In *Advances in Information Retrieval*, C. Gurrin, Y. He, G. Kazai, U. Kruschwitz, S. Little, T. Roelleke, S. Rüger, and K. van Rijsbergen, Eds., vol. 5993 of *LNCS*. Springer, 2010, 658–661.
- Clements, M., Serdyukov, P., de Vries, A. P., and Reinders, M. J. Using flickr geotags to predict user travel behaviour. In *Proc. SIGIR ’10*, ACM (2010), 851–852.
- De Choudhury, M., Feldman, M., Amer-Yahia, S., Golbandi, N., Lempel, R., and Yu, C. Automatic construction of travel itineraries using social breadcrumbs. In *Proc. HT ’10*, ACM (2010), 35–44.
- Girardin, F., Calabrese, F., Fiore, F. D., Ratti, C., and Blat, J. Digital footprinting: Uncovering tourists with user-generated content. *IEEE Pervasive Computing* 7 (2008), 36–43.
- Hägerstrand, T. What about people in Regional Science? *Papers in Regional Science* 24, 1 (1970), 6–21.
- Hassenzahl, M., Burmester, M., and Koller, F. AttrakDiff: Ein Fragebogen zur Messung wahrgenommener hedonischer und pragmatischer Qualität. *Mensch & Computer 2003. Interaktion in Bewegung* (2003), 187–196.
- Hecht, B., Schöning, J., Erickson, T., and Priedhorsky, R. Geographic human-computer interaction. In *Proc. CHI ’11 EA*, ACM (2011), 447–450.
- Hochmair, H. Towards a classification of route selection criteria for route planning tools. *Developments in Spatial Data Handling* (2004), 481–492.
- Kenteris, M., Gavalas, D., and Economou, D. Electronic mobile guides: a survey. *Personal and Ubiquitous Computing* 15 (2011), 97–111.
- Kjeldskov, J., Graham, C., Pedell, S., Vetere, F., Howard, S., Balbo, S., and Davies, J. Evaluating the usability of a mobile guide: the influence of location, participants and resources. *Behaviour and Information Technology* 24, 1 (2005), 51–65.
- Kostakos, V., Nicolai, T., Yoneki, E., O’Neill, E., Kenn, H., and Crowcroft, J. Understanding and measuring the urban pervasive infrastructure. *Personal Ubiquitous Comput.* 13, 5 (2009), 355–364.
- Kurashima, T., Iwata, T., Irie, G., and Fujimura, K. Travel route recommendation using geotags in photo sharing sites. In *Proc. CIKM ’10*, ACM (2010), 579–588.
- Lu, X., Wang, C., Yang, J.-M., Pang, Y., and Zhang, L. Photo2trip: generating travel routes from geo-tagged photos for trip planning. In *MM ’10*, ACM (2010), 143–152.
- Okuyama, K., and Yanai, K. A travel planning system based on travel trajectories extracted from a large number of geotagged photos on the web. In *PCM* (2011).
- Schöning, J., Hecht, B., and Starosielski, N. Evaluating automatically generated location-based stories for tourists. In *CHI EA ’08*, ACM (2008), 2937–2942.
- Shoval, N., and Isaacson, M. Sequence alignment as a method for human activity analysis in space and time. *Annals of the Association of American Geographers* 97, 2 (2007), 282–297.
- Snaveley, N., Garg, R., Seitz, S. M., and Szeliski, R. Finding paths through the world’s photos. In *SIGGRAPH ’08*, ACM (2008), 15:1–15:11.
- Suh, Y., Shin, C., Woo, W., Dow, S., and Macintyre, B. Enhancing and evaluating users’ social experience with a mobile phone guide applied to cultural heritage. *Personal Ubiquitous Comput.* 15 (August 2011), 649–665.
- Tuite, K., Snaveley, N., Hsiao, D.-y., Tabing, N., and Popovic, Z. Photocity: training experts at large-scale image acquisition through a competitive game. In *CHI ’11*, CHI ’11, ACM (2011), 1383–1392.
- Vermeulen, J. Improving intelligibility and control in ubicomp. In *Proc. Ubicomp’10*, ACM (2010), 485–488.
- Wilson, C. Activity patterns in space and time: calculating representative hagerstrand trajectories. *Transportation* 35 (2008), 485–499.
- Zhang, H., Law, E., Miller, R. C., Gajos, K. Z., Parkes, D. C., and Horvitz, E. Human computation tasks with global constraints. In *CHI’12*, ACM (2012), 217–226.