# Photowalking the city: Comparing hypotheses about urban photo trails on Flickr


Martin Becker
University of Würzburg
becker@informatik.uni-wuerzburg.de

Philipp Singer
GESIS - Leibniz Institute for the Social Sciences
Cologne, Germany
philipp.singer@gesis.org

Florian Lemmerich
University of Würzburg
lemmerich@informatik.uni-wuerzburg.de

Andreas Hotho
University of Würzburg
hotho@informatik.uni-wuerzburg.de

Denis Helic
Graz University of Technology
dhelic@tugraz.at

Markus Strohmaier
GESIS - Leibniz Institute for the Social Sciences and University of Koblenz-Landau
markus.strohmaier@gesis.org



## ABSTRACT

Understanding the way people move through urban areas represents an important problem that has implications for a range of societal challenges such as city planning and evolution, public transport or crime. In this paper, we lay our focus on a particular kind of human movement trajectories: urban photo trails, i.e., geo-temporal trails that are produced by humans when taking photos in urban areas. We study these movement trajectories and explore different explanations of how they materialize. For our experiments, we obtain trails of geo-temporally tagged photos from Flickr and adopt a Bayesian framework called HypTrails to study human movement. We specify a set of hypotheses and compare them on data obtained from four different cities (Berlin, London, Los Angeles, New York). Our results suggest that urban photo trails exhibit interesting commonalities and differences across cities that can be identified and explored through the approach adopted in this work.



**Categories and Subject Descriptors:** H.5.3 [**Information Interfaces and Presentation**]: Group and Organization Interfaces—*Web-based interaction*

**Keywords:** Human Trails; Flickr; Hypotheses; HypTrails


## 1. INTRODUCTION

Understanding the way people move through urban areas represents an important problem that has implications for a range of societal challenges such as city planning and evolution, public transport or crime. Recent research has studied human movement trajectories in cities through a variety of data sources including mobile phone data [15, 36], GPS tracking [42], WiFi tracking [29], location-based social media platforms [10], online photo sharing sites [12, 13, 14] and others. Such studies of human trails have deeply enriched our knowledge of the ways in which people move through urban space. For example, research has indicated that human mobility exhibits regularities [15, 36] and temporal and spatial patterns [10]. Research has also shown that we can successfully leverage these patterns. To give an example, De Choudhury et al. [12] illustrated that by looking at how people subsequently take photos in a city, high quality travel itineraries can be constructed automatically.

In this paper, we want to extend this stream of research by looking at trails of human movement exemplified in urban photo trails, and explore different explanations of how they have been produced.

**Problem and Objective.** In this work, we aim to *compare a set of different hypotheses about urban photo trails* given actual sequential photo data from the Web. We want to explore the plausibility of different potential explanations for the trails of photos that we can observe on social sharing websites such as Flickr. We define an *urban photo trail* as a sequence of spatial positions in a city over a period of time as, for example, obtained from the geo-temporal metadata of photos. Then, *hypotheses* can be expressed as different beliefs about transitions between spatial positions. For example, we might want to compare a *proximity hypothesis* —that represents a belief that humans frequently take subsequent photos in geographically close regions of a city— with a *Points-of-Interest (POI) hypothesis* that represents a belief that humans take subsequent photos of POIs. The problem is complicated by the fact that the photo trails produced by different user populations in different cities might yield different explanations given different data.

For a better understanding, consider our visualizations in Figure 1. In Figure 1(a) we depict Berlin and our cell based grid layout of the city. The heatmap illustrates the frequency of photos taken in each single cell of the grid as derived from our data. In this paper, we are interested in how urban photo trails are produced, i.e., how people move through the city while taking photos. Thus, we are interested in transitions and their probabilitites as the example in Figure 1(b) demonstrates. We see the transition probabilities that denote how likely people will take a photo in corresponding geographic cells next after they have taken one at the Bundestag. For in-

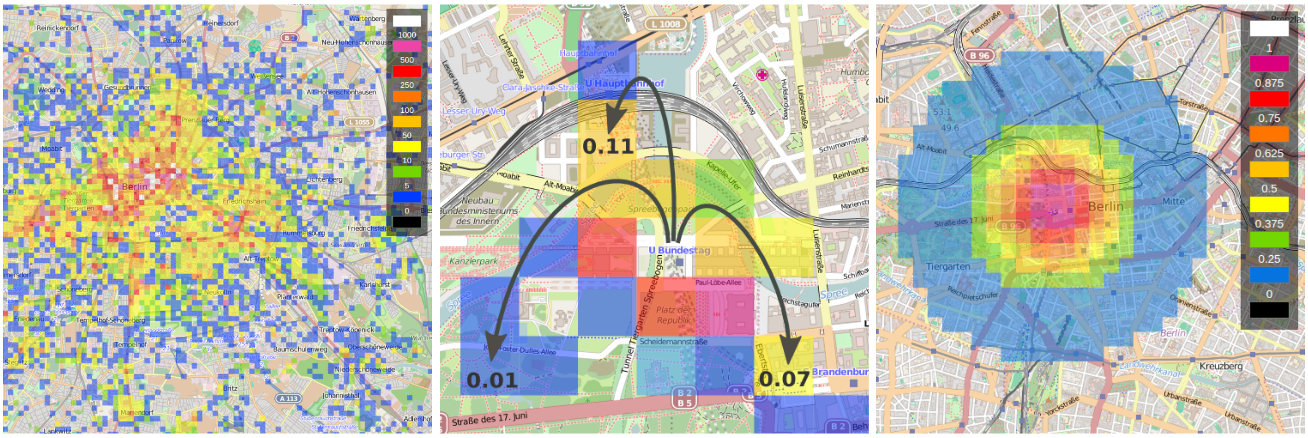


(a) Grid cells and their frequency (b) Transition probabilities (c) Proximity hypothesis

**Figure 1: *Main concepts.* This figure illustrates the main concepts of this work. In (a), we visualize the cell based grid layout of Berlin with corresponding frequencies of photos taken in corresponding cells as derived from our data at interest. In (b), we exemplary pick the cell (state of a Markov chain) where the German Bundestag is located, and visualize the transition probabilities of subsequent cells people frequently take photos at. For instance, with a probability of 0.07, people take a picture at the Brandenburg Gate after they have taken one at the Bundestag. Finally, in (c), we depict an exemplary proximity hypothesis that we want to study. This hypothesis believes that human successively take photos in proximate areas of a city. We express this as beliefs in Markov transitions. As (c) shows, this belief suggests that we have a higher probability that people take their next picture in direct proximity of the previous cell.**

stance, the probability of taking a photo at the Brandenburg Gate after taking one at the Bundestag is 0.07. As mentioned, the ultimate goal of our work is to compare hypotheses with each other. In Figure 1(c) we demonstrate the main idea of the exemplary proximity hypothesis mentioned above. The heatmap depicts the belief in transition probabilities for this hypothesis; higher proximity refers to higher belief.

**Approach.** To tackle these challenges, we resort to Bayesian approach called *HypTrails* [32] that is suited for comparing hypotheses about human trails. This approach is based on Markov chain modeling and Bayesian inference and allows for relative comparison of hypotheses given data. The main idea is to elicit priors from hypotheses and then utilize the sensitivity of the prior on the evidence (marginal likelihood) as a means to give insights into the relative plausibility of hypotheses. We reconstruct urban photo trails from geo-tagged user photo streams on Flickr. Then we construct hypotheses from general concepts like proximity, but also employ semantic web content from Wikipedia, DBpedia and YAGO to refine these hypotheses. The hypotheses are then compared on four different cities, namely Berlin, London, Los Angeles and New York.

**Contributions and Findings.** The overall contribution of this work is a systematic evaluation of hypotheses to explain how urban photo trails materialize in 4 different cities. We find that the cities share common characteristics, i.e., that the partial order of hypotheses for different cities is quite stable. Furthermore, we find that using semantic web content for hypothesis generation can contribute to explaining the production of the photo trails of humans at interest. Concretely, the general concept of proximity combined with Wikipedia entities in combination with their page views yields the highest evidence among the hypotheses studied.

**Structure.** We start by giving an overview of related work in Section 2. Next, we describe the utilized HypTrails method for comparing hypotheses about human photo trails as well as the Flickr data studied in Section 3. In Section 4 we elaborate the hypotheses at interest and how we can express them. Afterwards, Section 5 shows our experiments and results. We discuss our work in Section 6 and conclude it in Section 7.

## 2. RELATED WORK

This section is structured to cover related work from three areas of research: (i) studies on geo-spatial trails, (ii) studies on Flickr as well as (ii) studies on human trails on the Web in general.

**Geo-spatial trails.** In the past, geo-spatial trails have frequently been studied by looking at mobility patterns of humans as derived from mobile data. Song et al. [36] studied the limits of predictability in these kind of geo-spatial trails. They found high predictability with a lack of variability between humans. This suggests that regularities emerge when people move through the world. The authors even argue that these regularities might be based on the inherent regularities in human behavior in general. [15] suggests similar observations by looking at mobile phone data: human mobility trails show high temporal and spatial regularity. In [10], the authors investigated basic laws that steer human mobility as captured from location-based social media platforms and mobile phone data. Even though these kind of datasets differ, the results demonstrate several common patterns such as periodicity and influence by social ties. However, short-ranged human mobility does not seem to be impacted by the social network structure at the same level as long-distance travel. Previous research has also focused on studying the network properties of geo-spatial trails [23].

**Flickr.** Geo-spatial trails have also been studied on Flickr. For instance, De Choudhury et al. [12] aimed at leveraging photo trails for automatically constructing travel itineraries through cities. The main idea of their approach is to aggregate all trails into a POI graph and then use the popularity of POIs for calculating itineraries. The authors' evaluation based on a comparison against well-known itineraries from travel agencies indicates high quality. Similar work has been done by Tai et al. [37] who have used past landmarks photographed by users for recommending sequences of new landmarks as derived from sequential information by other users on Flickr. Girardin and colleagues have conducted several studies on photo trails as captured from Flickr. In [13] they studied digital footprints as determined explicitly via photo trails and in [14] they focused on tourist dynamics based on concentrations and spatio-temporal flows revealing popular points of interests, density points and common trails tourist follow.

Apart from trails, Flickr has also been studied in several other dimensions, specifically regarding tagging. Tagging research on Flickr has been steered by the early work by Marlow et al. [20]. As tags play a crucial role on Flickr, related works have also focused on recommending tags on Flickr such as [31]. Tags can even reveal the gender and location of users [28]. Apart from tags, photos themselves can give insights into visit times at landmarks [27] or for automatically finding places that people find interesting to photograph [11]. Lerman and Jones [18] demonstrated that browsing through the contacts' photo streams is a primary method for Flickr users to find images. Additionally, related works have also studied the social network on Flickr [7, 22].

**Human trails on the Web.** Human trails—i.e., sequential interactions of humans with the Web—have been studied by our research community in various contexts since the inception of the World Wide Web [2]. While some work has mainly focused on modeling (e.g., [3, 4, 9, 26, 30, 33]), others have been more interested in investigating regularities, patterns and strategies (e.g., [16, 38, 39]) that emerge when humans sequentially engage with the Web. These kind of works have been interested in answering similar objectives as this paper, i.e., understanding sequential steps by humans. In the following, we want to provide some exemplary related works in that direction.

One of the most prominent interaction of humans with the Web—maybe even also one of the first—is the navigation between websites. An example of early work regarding patterns and strategies in human navigational trails is by Catledge and Pitkow [6] who aimed at augmenting WWW pages. Further work [9, 16] emphasized that humans base their navigational steps on some regularities and patterns. As an example, Pirolli and Card [25] derived the so-called *information foraging theory* that postulates that human behavior in an information environment on the Web is guided by *information scent* which is based on the cost and value of information with respect to the goal of the user [9]. Also, researcher have found that semantics influence human navigational choices [5, 8, 24, 34, 39, 40].

Apart from human navigational trails, previous work has also identified cognitive strategies in other kinds of human trails on the Web. For example, An et al. [1] found that human participate in partisan sharing on Facebook. White and Huang [41] identified the importance of following search trails on the Web. The work by Yang et al. [43] states that humans follow certain stages in their sequential behavior which the algorithm proposed in corresponding work can identify. The work by Matsubara et al. [21] emphasizes the existence of trends in human trails that can be captured.

Contrary to elaborated previous works, we directly aim at comparing hypotheses about human trails on the Web in this article. This means that we do not strive for finding explanatory patterns, but rather we want to directly judge about the relative order of hypotheses that aim at explaining the sequential behavior. To that end, we apply the so-called HypTrails approach as presented in [32], as introduced in some more detail in the following section.

# 3. METHODS AND MATERIAL

In this section, we give an overview of our methodology and the data that we study in this article. By and large, we are interested in studying trails that emerge when humans take photos in a geographic environment. In detail, we want to get insights into the relative plausibility of hypotheses about how humans choose their next location to take a photo by comparing them. For doing so, we resort to an approach called HypTrails as well as data derived from Flickr. We describe both the method and the data next.

## 3.1 Methodology

HypTrails is an approach for comparing hypotheses about human trails with each other. Technically, HypTrails is based on Markov chain modeling and Bayesian inference. The approach provides insights into the relative plausibility of hypotheses—assumptions in common and uncommon transitions. In the following, we only very shortly outline the main concepts and ideas of HypTrails and refer the reader to [32] for more detailed information.

With HypTrails, we model trails as a first-order Markov chain—a stochastic system that models transitions between states. In Bayesian statistics, the marginal likelihood—also called *evidence*—denotes the probability of the data given a hypothesis $H$. The main idea of HypTrails is to utilize the influence of the prior on the evidence for comparing hypotheses with each other. In particular, hypotheses are expressed as Dirichlet priors. Thus, using different hypotheses as priors leads to different marginal likelihoods when combined with observed data. When we compare two hypotheses, a higher evidence for a given hypothesis indicates a higher plausibility of it. Bayes factors are utilized for determining the strength of evidence for one hypothesis over the other. In this article, all Bayes factors are decisive which is why we refrain from explicitly reporting them—please refer to [32] for further details.

In detail, the following steps are necessary for an application of HypTrails given a set of generic hypotheses about the production of the human photo trails at interest:

(i) First, we need to specify a hypothesis matrix $Q$ for each hypothesis. $Q$ quantifies our assumptions about transitions between the states observed in the trails. No negative values are allowed and higher values correspond to stronger assumptions. We describe this process as well as our hypotheses at interest in detail in Section 4.

(ii) Next, we need to pass these matrices and observed data (see Section 3.2) to HypTrails that subsequently elicits the Dirichlet priors for each hypothesis with varying values of the hypothesis weighting factor $K$. Basically, the higher we set the parameter $K$ the stronger we believe in a given hypothesis.

(iii) Based on this elicitation, HypTrails determines the evidences for each hypothesis and each parameter $K$. As mentioned above, for simplicity, we can state that one hypothesis

$H_1$ is more plausible compared to another hypothesis $H_2$, if the evidence of $H_1$ is higher than the one of $H_2$ for the same value of $K$. Thus, the partial ordering based on the plausibility of respective hypotheses $\mathbf{H} = \{H_1, H_2, ..., H_n\}$ can be determined by ranking their evidences from largest to smallest for single values of $K$. We present corresponding results in Section 5

## 3.2 Datasets

Throughout this work, we are interested in studying human photo trails through cities by analyzing Flickr[1] data. In this section, we first describe the dataset generation process, that is, what data was collected and how we transform it into the required representation of trails and state transitions. Then, we highlight some basic characteristics of our datasets.

**Data collection.** Our datasets[2] contain metadata—i.e., user, temporal and geo-reference (latitude and longitude) data—about images uploaded to the Flickr platform. In particular, we focus on pictures taken in the cities of Berlin, London, Los Angeles and New York between January 2010 and December 2014. For each city, we define a bounding box, see Table 3.2. We acquired corresponding data by crawling Flickr's public API. Since our analysis requires an exact position, we additionally filter the pictures by the positional accuracy as indicated by the Flickr API, concentrating on pictures with street-level accuracy (level 16 on the Flickr scale[3]) only.

For our analysis, we interpret the sequence of all images of a single user, ordered by the time the photo was taken, as a trail regardless of the time difference between pictures. Each node in a trail, i.e., each image, is mapped to a cell of a grid that we place over the respective city according to its geo-reference . The grid cells are used as states of the Markov chain in the HypTrails approach. For our experiments, we choose a size of 200m x 200m for each cell. Figure 1 show several grids and grid cells in Berlin to give an idea about the chosen granularity.

In all of our datasets, the trails contain sequences of pictures taken in the same grid cell. However, since we are mostly interested in the sequence of different places humans visit and photograph, we focus on the sequence of visited places in this paper. For that purpose, we remove all but one subsequent occurrences of the same grid cell in the paths. In other words, all transitions in the Markov chain from a state to itself are removed. In the end, we work with trails over cells in a city where each trail corresponds to the places one person has successively taken pictures at.

[1]https://flickr.com

[2]Dataset access can be requested via e-mail.

[3]see https://www.flickr.com/services/api/flickr.places.findByLatLon.html.

**Table 1: Bounding boxes and center coordinates used for data collection and hypothesis creation.**

| | min lon. | min lat. | max lon. | max lat. |
|---|---|---|---|---|
| Berlin | 13.088400 | 52.338120 | 13.76134 | 52.675499 |
| London | -0.5103 | 51.2868 | 0.3340 | 51.6923 |
| Los Angeles | -118.6682 | 33.7037 | -118.1552 | 34.3368 |
| New York | -74.2589 | 40.4774 | -73.7004 | 40.9176 |

| | Berlin | London | Los Angeles | New York |
|---|---|---|---|---|
| lon | 13.383333 | -0.1280 | -118.2450 | -74.0071 |
| lat | 52.516667 | 51.5077 | 34.0535 | 40.7146 |

On overview of the filtered processed dataset can be seen in Table 3.2. The number of photos corresponds to the number of trail nodes, i.e. the number of photos after filtering self-transitions and trail of length smaller than 2. The number of cells is the number defined by the corresponding bounding box and grid. Covered cells is the number of cells covered by the trails.

**Points of interest.** An intuitive hypothesis (see Section 4) for photo trails is a hypothesis where pictures are taken at important places within a city, so called points of interest (POIs). To identify POIs, we queried linked data from the DBpedia [17] and YAGO projects [19]. In particular, we extracted all entities from DBpedia that are located in a city's bounding box according to the properties *geo:lat* and *geo:long*. The result set of entities also includes areas such as "Germany", which are also mapped to a specific location in the bounding box according to these properties. In a data cleaning step, we removed all entities that are specified as districts (i.e., having the value *yago:District108552138* as *rdf:type*).

Additionally, we quantify importance of a POI in some hypotheses. As an approximate measure of importance we take pageview counts of the Wikipedia articles describing the POIs. We extract view counts from the Wikimedia download page[4]. In this work, we use the view counts for January 2012. Table 3.2 shows the number of POIs per city and their average view count.

# 4. HYPOTHESES

The HypTrails approach allows to compare hypotheses about human trails—see Section 3 for details. These are easily specified by constructing matrices that reflect beliefs about transitions from one state to another. In our setting, that means, that we hypothesize where the next picture of a person will be taken given only the location of her last picture. In this section, we discuss how intuitions and beliefs about the production of human photo trails can be formalized into hypothesis matrices according to the concepts of the HypTrail approach.

A single hypothesis matrix $Q$ is defined as follows: for each geographic grid cell (state in the Markov chain) $s_i$ we set our beliefs about transitions to every other grid cell $s_j$. In the HypTrails approach, see [32] and Section 3, higher values should correspond to higher beliefs in corresponding transitions. In this paper, we express our assumptions for a given hypothesis directly as row probabilities $P(s = s_j|s_i)$. To give an example, one hypothesis might assume that people choose state (grid cell) $s_2$ with a probability of 0.5 for taking their next photo when they previously have taken a photo in cell $s_1$. Thus, this hypothesis would set $P(s = s_2|s_1) = 0.5$. Also

[4]http://dumps.wikimedia.org/other/pagecounts-raw/

**Table 2: Basic dataset statistics.**

| | Berlin | London | Los Angeles | New York |
|---|---|---|---|---|
| years | 2010-11 | 2010-14 | 2010-14 | 2010-14 |
| photos | 60978 | 794535 | 300373 | 714549 |
| cells | 43052 | 66444 | 84014 | 58065 |
| trails | 4364 | 35101 | 15357 | 31246 |
| covered cells | 6343 | 23694 | 25834 | 15232 |
| avg. trail length | 13.97 | 22.64 | 19.56 | 22.87 |
| | | | | |
| POIs | 1085 | 7228 | 1462 | 6002 |
| avg. view counts | 1240 | 1272 | 3654 | 1511 |

consider our example depicted in Figure 1(c). Ultimately, all elements of the hypothesis matrix Q are set according to this specification: $Q(i,j) = P(s = s_j|s_i)$. With this notation, each row $i$ of $Q$ has $\ell_1$-norm leading to $\sum_j Q(i,j) = 1$.

In most cases, we will not directly define a belief as a probability of each transition $P(s = s_j|s_i)$, but as a function $\bar{P}(s = s_j|s_i)$. This function can then be transformed into a probability distribution by a normalization factor $\frac{1}{Z}$ obtained by summing over all values of $\bar{P}$:

$$P(s = s_j|s_i) = \frac{1}{Z}\bar{P}(s = s_j|s_i),\ Z = \sum_{k=1}^{n} \bar{P}(s = s_j|s_i)$$

In general, we can categorize hypotheses into two types: for *global hypotheses* our belief in transitions to the next state are independent from the last state. In this case, the rows of the hypothesis matrix are all identical to each other. By contrast, for *local hypotheses* beliefs in transitions to the next state do depend on the previous state. In the following, we describe all hypotheses studied in this article and how we can transform them to hypothesis matrices suitable for the HypTrails approach.

## 4.1 Uniform Hypothesis

This global hypothesis formulates the belief that each transition from any grid cell to another is equally likely. In other words, it assumes that users will randomly take pictures anywhere in the city regardless of the previous location. This is formalized as follows:

$$P_{uni}(i,j) = \frac{1}{n}$$

Where $n$ is the number of grid cells in the defined bounding box—i.e., the number of states in the state space. This hypothesis can be seen as a baseline; if other informed hypotheses are not more plausible than this one, we cannot expect them to express good hypothesis about the production of human photo trails.

## 4.2 Center Hypotheses

Typically, the center of a city is its most lively part and one could expect that most pictures are taken there. Therefore, this global hypothesis believes that the next picture of a user is always taken with a higher probability near the city center, regardless of the position of the user's last image. This is formalized by computing the geographic center $C$ of the city (in our case given by the center of the cities bounding box) and employing a two-dimensional Gaussian distribution centered at this point. Given the geographic (haversine [35]) distance $dist(C,i)$ between the city center $C$ and the central point of the grid cell $s_i$, the entries of the hypotheses matrices can be derived from the following distribution:

$$\bar{P}_{center}(i,j) = e^{-\frac{1}{2\sigma^2} dist(C,j)^2}$$

Note, that in order to obtain a proper probability mass function a normalization constant has to be applied since the number of states in our model is finite.

The hypotheses are parametrized by specifying the standard deviation $\sigma$ (e.g., in km) of the Gaussian distribution. A small value of $\sigma$ indicates assumes that most pictures are taken very close to the city centre. If $\sigma$ approaches infinity, the hypotheses approximates the uniform hypothesis.

## 4.3 Proximity Hypotheses

This hypothesis expresses the belief that the next image of a user will be taken nearby the last image considering the geographic distance. This hypothesis is partly motivated by findings of previous work [10, 15, 32]. We formalize this as follows: we define the distance $dist(i,j)$ between two grid cells $s_i$ and $s_j$ as the distance between the geo-coordinates of the centers of the two cells. Then, we can specify the respective believed transition probabilities by employing two-dimensional Gaussian distributions.

$$\bar{P}_{prox}(i,j) = e^{-\frac{1}{2\sigma^2} dist(i,j)^2}$$

As before, a standard deviation $\sigma$ must be specified. Here, a small value of $\sigma$ suggests a picture is more likely to be taken very close to a user's previous picture. An example for this hypothesis is depicted in Figure 1(c) where we depict our beliefs in transitions from one state to other states (i.e., one row of $Q$).

## 4.4 Points of Interest Hypotheses

Previous work has shown that photo trails as derived from Flickr can be utilized to automatically construct travel itineraries through a city [12]. This suggests that humans move along points of interests (POIs) when taking photos throughout urban spaces. Examples of points of interest are not only tourist attractions, but also stations for public transportation or the locations of governmental institutions. For our work, we extracted potential points of interest from DBpedia, see Section 3.2. Thus, we are interested in studying these global hypotheses that expresses the belief that people take a majority of pictures near such points of interest. However, it is save to assume that pictures are often not taken directly where a POI is located in the data, either as the POI itself covers a somewhat larger area, or as pictures have to be taken at a certain distance from the POI (e.g., to cover a architectural attraction fully in one picture), or as people intend to visit the POI, but find something interesting to photograph nearby. Again, we model this by using a multivariate Gaussian distribution, hypothesizing that there is a high chance that a picture is taken directly at the POI and a lower chance that a picture is taken in a grid cell in the neighborhood of a POI. Using this normal distribution, we get for each cell $c$ and each POI $q \in Q$ a probability $P_{poi}(q,c)$ that $q$ generates a picture in the cell $c$:

$$G(q,c) = e^{-\frac{1}{2\sigma^2} dist(q,c)^2},$$

as before $dist(q,c)$ describes the haversine distance between the POI $q$ and the cell $c$. Then, for each cell, we aggregate the probabilities of all points of interests:

$$\bar{P}_{poi}(i,j) = \sum_{q \in Q} G(q,j)$$

Similar to the distance hypothesis, the standard variation $\sigma$ has to be chosen. In this case, a small parameter indicates that one assumes pictures to be taken directly at the point of interest, a larger setting assumes that pictures are taken somewhere in the surrounding of a POI.

## 4.5 Weighted Points of Interest Hypotheses

Each city contains a large amount of potential points of interest. However, not all might be of equal importance. For example, the "Brandenburg Gate" could be expected to have more influence on human trails in the city of Berlin than the less known "Charlottenburg Gate". Thus, for this hypothesis, we assume that the importance of a POI $q$ can be approximated by the view count $views(q)$ of the Wikipedia article corresponding to the POI. If the view count of an article is very high (as e.g., for the "Brandenburg Gate"), we expect the respective POI to have a stronger influence on the sequence of image locations. We formalize this by introducing weighting functions to each summand of the POI hypothesis, resulting in:

$$\bar{P}_{weighted\_poi}(i,j) = \sum_{q \in Q} (views(q) \cdot G(q,j)) .$$

This linear weighting potentially overemphasizes the importance of large points of interest. Therefore, we suggest as variation to use logarithms of view counts as weights instead:

$$\bar{P}_{log\_weighted\_poi}(i,j) = \sum_{q \in Q} (\log(views(q)) \cdot G(q,j)) .$$

## 4.6 Combinations of Hypotheses

Some proposed hypotheses are to a large degree independent from each other and can easily be combined. Formally, we combine hypotheses by multiplying the respective elements of the hypothesis matrices. We study two combinations throughout this paper, but other kinds of combinations are conceivable.

**Proximate weighted points of interest hypotheses.** As an example, an intuitive hypothesis is that a user will take her next picture at some POI, but more likely at a POI that is close to the position of her last picture. We can express such a hypothesis by combinations of distance hypotheses and POI hypotheses can be computed as

$$\bar{P}_{prox\_poi}(i,j) = \bar{P}_{prox}(i,j) \cdot \bar{P}_{weighted\_poi}(i,j.)$$

**Proximate center hypotheses.** Similarly, hypotheses that express the belief that the next picture is taken more likely near the position of the last image, but also more likely near the city center can be determined by:

$$\bar{P}_{prox\_center}(i,j) = \bar{P}_{prox}(i,j) \cdot \bar{P}_{center}(i,j.)$$

# 5. EXPERIMENTS

In Section 4, we have introduced a set of hypotheses which provide possible explanations for human trails while moving through a city and taking pictures. In this section, we compare these hypotheses with each other based on empirical trails derived from four different cities, i.e., Berlin (Germany), Los Angeles (USA) and London (United Kingdom) (see Section 3.2). For doing so, we resort to our methodological approch outlines in Section 3.1. Since each hypothesis has different parameters and some hypotheses depend on one another, we first conduct an in-depth study of each hypothesis on Berlin in the same order as they have been introduced in Section 4 and compare the results. Afterwards, we report the best hypotheses (based on their parameter configurations) of each city and highlight prominent differences between cities.

In the following, we are often referring to Gaussian distributions used for weighting transition probabilities or factors. In this context, the elements of the hypothesis matrix $Q$ might be rather small. Thus, if not stated otherwise, we always set the value for a belief in a transition $Q(i,j) = \bar{P}(s = s_j|s_i)$ to 0 if $\bar{P}(s = s_j|s_i) < 0.01$ (cf. Section 4). Also, as stated in Section 3.2, since we are interested in how people move around the city and not if they stay and how long they stay at a certain spot, we have removed self transitions from trails and consequently set $Q(i,j)$ to 0 for all self-transitions in every hypothesis.

## 5.1 Berlin

In this section, we introduce each hypothesis in detail using the example of Berlin in the same order as they have been introduced in Section 4.

**Uniform Hypothesis.** The first hypothesis introduced was the uniform one (see Section 4). It represents the belief that the next picture will be taken anywhere within the specified grid and cell layout of a city, no matter where a person has previously taken a picture at. It can be seen as a baseline we want to compare other hypotheses against. If another hypothesis is not more plausible than the uniform one, we cannot expect it to provide a good explanation of how people move around a city while taking photos.

**Center Hypotheses.** With these hypotheses, we hypothesize that the daily activities of the local population of a city and especially of the largest part of the visiting tourists mainly focus on the center of the city. For Berlin, the main part of the photos is clearly centered around the cultural center as can be seen in Figure 1(c). Consequently, we expect the center hypothesis, i.e., the belief that people move towards the center and stay there while taking pictures, to be a better explanation of human photowalking behavior compared to the uniform hypothesis. We calculate the evidence of the center hypothesis around the city center for four different spreading parameters, i.e., standard deviations of a Gaussian distribution: 1km, 3km, 5km and 10km. For the center coordinates used refer to Table 3.2. The results can be seen in Figure 2(a). Throughout all values of $K > 0$, we observe the highest plausibility for a standard deviation of 3km indicating a rather centralized city center since standard deviations of 5km and 10km result in clearly less evidence. The standard deviation of 1km yields good results for small K compared to 5 and 10km sigmas. The larger we set K, the stronger we believe in a given hypothesis. Thus, fewer specific parameter configurations receive higher prior probability tailored towards those represented by the hypothesis—please see [32] for further details. This means that the higher we set K, the importance of the center peak with a given sigma becomes more prominent. For instance, for a center hypothesis with a sigma of 1km, lower values of K might also have (lower) beliefs in parameter configurations outside this radius. Based on this, we can explain the falling evidence for higher values of K for the center hypothesis with a 1km sigma value. This hypothesis seems to concentrate too densely on the center with a small radius without also considering further surroundings of the center. Thus, higher values of K increase the belief in this density, why a falling evidence can be observed. Contrary, other values of sigma seem to capture the radius better, resulting in an increase of their evidence with higher values of K. As mentioned, a radius of 3km results in the highest evidences. As expected, all center hypotheses are more plausible compared to the uniform one.

**Proximity Hypotheses.** The proximity hypothesis as introduced in Section 4 is grounded on the belief that people are mostly walking around when taking pictures in a city. Consequently, pictures will be taken in close proximity to the

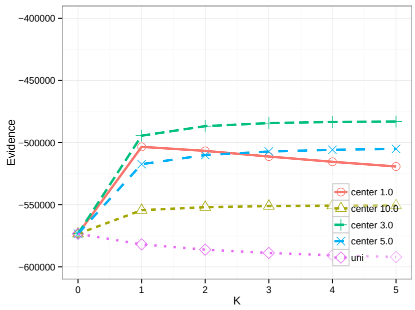


(a) Center hypotheses

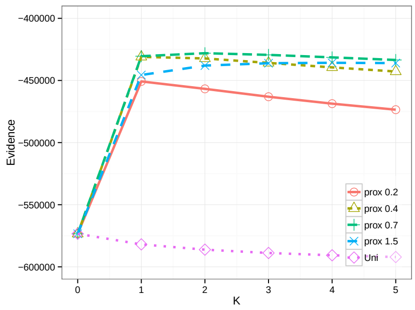


(b) Proximity hypotheses

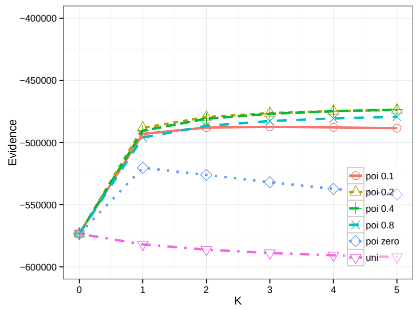


(c) POI hypotheses

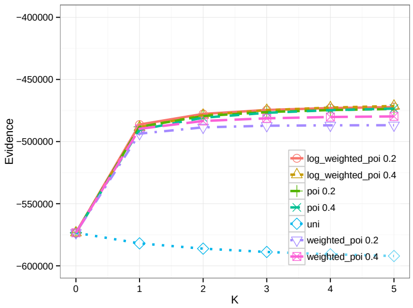


(d) Weighted POI hypotheses

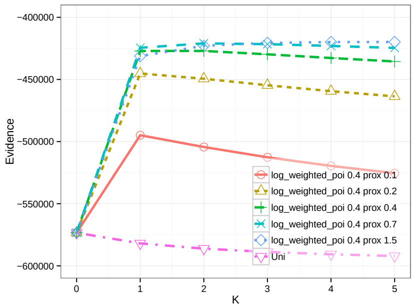


(e) Proximate weighted POI hypotheses

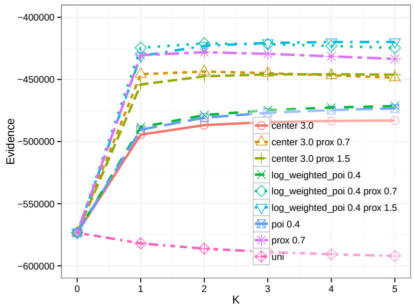


(f) Comparison between best hypotheses

**Figure 2: *Berlin experiments.* This figure visualizes the results for our studies on human photo trails derived from Berlin. First, for each type of hypotheses at interest, we compare the plausibility of various parameter configurations with each other (a-e). Then, in (f) we compare the best hypotheses of each set with each other. The results reveal intriguing explanations for human movement behavior while taking photos in Berlin. First of all, all hypotheses at interest are more plausible than a uniform hypothesis that can be seen as a baseline. The various parameter configurations reveal that the level of proximity is important. Being to narrow or being too ample is worse compared to covering the right proximity. For instance, in (b) we can see that a radius of 700m lads to the most plausible proximity hypothesis. Overall (f), a combination of the proximity hypothesis and the weighted POI hypothesis provides the best explanation. This suggests that people prefer to subsequently take photos at proximate POIs in a city. For further details, please refer to Section 5.1.**

last picture. This idea is modelled by a Gaussian distribution around the current location. The farther away another cell is, the less likely it becomes that a picture is taken there next. We calculate the proximity hypothesis for four different standard deviations: 200m, 400m, 700m and 1.5km. The results are depicted in Figure 2(b). Again, this hypothesis is strongly superior to the uniform hypothesis. Comparing the different spreading parameters, a sigma of 700m produces the highest evidence for all values of $K > 0$. This indicates that the best proximity hypothesis is the one that assumes that follow-up pictures are being taken in a 700m radius from their predecessor. For spreading parameters of 200 and 400m a similar situation occurs as for the center hypotheses with a sigma of 1km: They seem to concentrate their belief on a too narrow proximity leading to decreasing evidence values for higher values of $K$. Thus, for $K > 3$, the proximity hypothesis with a spreading factor of 1.5km also overtakes the one with a sigma of 700m.

**Points of Interest Hypotheses.** Since not every point in a city is equally popular, we formulate the POI (points of interest) hypothesis. It is similar to the center hypothesis in that regard, that the distribution is the same for every starting place. The difference is that instead of moving to the center of the city people, it expresses the belief that people move to any cell containing a POI; thus, avoiding places which do not contain a POI. The most simple way of weighting places by their importance is by taking the number of POIs they contain. Yet, this focuses the transition destinations to a very limited amount of cells. Thus, we employ a similar approach to the center hypothesis, letting each POI spread its influence to surrounding cells weighted by a Gaussian distribution based on its distance to the corresponding cell and sum the resulting influence of each POI at each cell. We extracted the POIs from Wikipedia and Yago as explained in Section 3. Again, we compare different spreading parameters: 100m, 200m, 400m and 800m. Figure 2(c) shows that the POI hypothesis without Gaussian spread indeed performs inferior to those POI hypotheses allowing their influence to spread. The two rather close-ranged spreads 200m and 400m perform the best implying that people indeed move towards POIs. The worse performance of too narrow and too wide ranges is an indicator that people tend to visit places and take pictures of the place at a close range, but not necessarily from the inside. Considering tourists, this might be an indicator of them documenting where they were: Close range and too far away pictures might not show enough detail. Overall the POI hypothesis seems to perform worse than the distance hypothesis comparing scales. Note that each POI is spreading a base value of 1 in this case. This limitation is lifted next, where we allow each POI to be associated with a weight.

**Weighted Points of Interest Hypotheses.** In contrast to the normal POI hypothesis, the weighted POI hypothesis allows POIs to exhibit different levels of importance when spreading their influence. The idea is that some POIs attract a greater number of people, thus, we might believe that the probability of a transition to such POIs is higher. In order to define importance levels for the defined POIs, we extracted online usage statistics of the corresponding Wikipedia pages. Specifically, we are using view counts to weight the POIs as described in Section 4 and 3. We are considering two schemes of deriving weights from the view counts. One is to use the view counts themselves and the other is taking the logarithm of the view counts. We do this to account for Fitt's law, i.e., that view counts are power law distributed, smoothing strong peaks of highly popular places like the Brandenburg Gate. Now, instead of calculating the weighted POI hypothesis for every spreading parameter, we chose the best of the normal POI hypothesis. In this case, two distribution parameters are tied, that is, 200m and 400m. Figure 2(d) shows the results. We observe that the raw view counts result in a lower evidence strengthening the need for a normalization of the view counts. Furthermore, we can see that using the view counts strengthens the evidence of the hypothesis even if marginal. This is an indicator that content and popularity of social media content is indeed in some way connected to human behavior in their daily live.

**Proximate Weighted Points of Interest Hypotheses.** Note that the POI hypotheses as well as the center hypotheses have both been neglecting local transition probabilities, i.e., the transition probabilities to the other cells are independent of the the cell the user starts from. As described in Section 4, we can introduce local sensitivity by combining hypotheses with the proximity hypothesis. We choose the weighted POI hypothesis for a sigma parameter of 400m since it was one of the best performing hypotheses so far and incorporates human behavior data. For reasons of space limitations, we choose the weighted POI hypothesis using logarithmic weights together with a spreading parameter of 400m to be combined with the proximity hypothesis. In addition to the weighting parameter used to build the weighted POI hypothesis, the proximity hypothesis also has a spreading parameter. We choose to evaluate the following: 100m, 200m, 400m, 700m and 1.5 km. Figure 2(e) shows the results. Compared to the other hypotheses the evidences are as high as for the proximity hypothesis with large spreading parameters being the most favorable. The results are best for a sigma of 700m and 1.5km. A direct comparison with the other hypotheses follows in the next section.

**Comparison.** For a direct comparison of the different hypotheses we are taking the most plausible ones (best parameters) of each set. This is the center hypothesis for 3km, the proximity hypothesis for 700m, the POI hypothesis for a spreading parameter of 400m, the log weighted POI hypothesis for 400m as well as the proximate, log weighted POI hypothesis for 400m and a spread of 700m and 1.5km. The main result as shown in Figure 2(f) is, that indeed, the combination of the proximity hypothesis and the weighted POI hypothesis provides the best explanation of how people move around Berlin while taking photos. Yet, we observe that the center hypothesis performs similar to the POI based hypotheses without the proximity factor. The reason might be that the POI hypotheses simply behave the same as the center hypotheses in a way that they believe in transitions to the city center since there are generally more POIs there. To contrast this line of thought, we include a proximate center hypothesis. If the former was true, this new combination should perform equally well. However, both variants are still strongly inferior to the proximate, weighted POI hypothesis. Thus, we can improve the proximity hypothesis by adding human generated content from Wikipedia and corresponding user statistics. This allows us to conclude that human trances on the Web and human traces in the real world are conceptually intertwined by shared entities such as sights being visited in real life and documented digitally.

## 5.2 Los Angeles, London and New York

To further augment the results from Section 5.1, we analyze three more cities, namely, Los Angeles (USA), London (United Kingdom) and New York City (USA). We show similarities and especially highlight some differences between the cities. To this end, we provide Figure 3 in addition to the results from Figure 2. Corresponding sub-figures are built in the same way as Figure 2. That is, we first determine the most plausible parameter configuration for each type of hypothesis, and then compare these in Figure 3(a), Figure 3(b) and Figure 3(c). However, due to space limitations, we resort from explicitly presenting the individual results and focus on the comparison here. By comparing the four different cities, for most parts, all hypotheses perform pretty much the same and the best parameters stay consistent. This indicates the hypotheses about picture trails in Berlin can be generalized to other cities as well implying some basic patterns that hold even across countries.

However, there are two exceptions which are worth mentioning. First, in Los Angeles (see Figure 3(a)) the most plausible center hypothesis is the one with a spread of 10km instead of 3km. This indicates a city that either has a very large center or non at all. Arguably, Los Angeles is a spread out city which may cause this divergence This is further confirmed by the fact that LA supports higher spreads for POI hypotheses, i.e., 400km, compared to London or New York City. Also, the the best performing hypotheses are strongly indicating decreasing evidence with increasing K. This is a strong indicator that the corresponding hypotheses are missing specific parameter configurations (empirical transitions), thus, not covering cells that should be modelled.

Second, the linearly weighted POI hypothesis in London is superior to the logarithmically weighted one. This may be due to differently distributed view counts and has to be further investigated in the future.

# 6. DISCUSSION

In this work, we have conducted rigorous experiments to gain a better understanding of the underlying processes that are employed when people take photos while moving through cities. We want to dedicate this section to discuss some potential limitations and future aspects that we see as worthy additions to the studies laid out in this work.

**Data restrictions.** While we have made efforts in comparing the trails of various large cities, our experiments still exhibit certain data restrictions. In the following, we want to shortly discuss two potential restrictions. (i) First of all, we limit the studies to human photo trails as derived from Flickr. Thus, we are only able to make judgements about behavioral aspects that emerge when people move through a city and take photos. We cannot be sure that our observations would be similar if we would study other kinds of movement data such as mobile phone data. Even though this not the goal of this article, we might assume that certain behavioral aspects

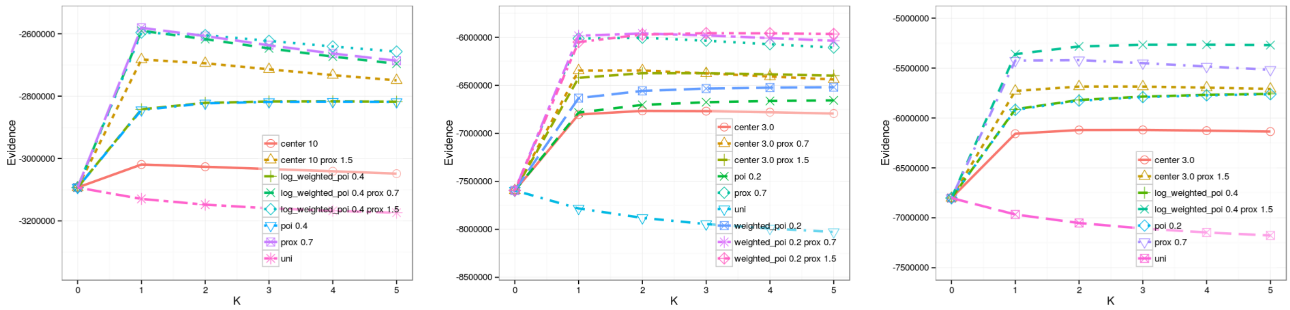


(a) Los Angeles (b) London (c) New York City

**Figure 3: *City (LA, LDN, NYC) experiments.* This figure visualizes the results for our studies on human photo trails derived from Los Angeles (a), London (b) and New York City (c). We only present a comparison of the best hypotheses for each type of hypotheses for each city. Overall, we can identify similar explanations across cities (cf. Figure 2(f)). One main exception is that we observe higher spread parameters for Los Angeles which might be explained by the fact that Los Angeles is a rather large and spread out city. Overall, our results show that our hypothesis that people prefer to subsequently take photos at proximate POIs in a city provides a good explanation for our human photo trails at interest across various cities.**

are similar, regardless which data we look at as suggested in [10]. Yet, we plan on delving deeper into this in future, by not only contrasting different cities, but also different kinds of movement data—e.g., social check-in data, mobile phone data or business reviewing data. (ii) Second, we focus on a city level in this work. While this allows us to gain insights into urban behavior, we might observe different behavior if we would extend our scope of interest. To give an example, by looking at cities, we already constrain our studies on a small geographic space which might benefit proximity based hypotheses to a certain extend. If we would extend the scope, by e.g., looking at a country or continent level, the results might largely differ. However, then, other kinds of hypotheses might be more plausible to study.

**Choice of hypotheses.** We focus on our experiments on a set of hypotheses. These are partly motivated by related work and partly by our own intuitions. While observed results give interesting insights into the production of the trails at interest, our set of studied hypotheses is not complete. All other kinds of hypotheses are conceivable and can be investigated with HypTrails and our data. Next, we suggest some potential candidates: (i) A river hypothesis might express the belief that a river is a natural barrier in a city. People might have restrictions between crossing a river in order to take pictures at the other side of it. (ii) A district hypothesis might have similar intentions, but in this case the district boundaries may be some kind of natural barrier. Finally, (iii) a demographic hypothesis might make certain assumptions about the influence of demographic aspects in a city. For example, people might prefer to take photos in parts of a city with less crime rates.

**Tourists vs. residents.** Previous work has suggested that the photographing behavior on Flickr differs between tourists and residents in a city [12, 14]. The authors of [12] argue that residents are not under the direct pressure of visiting as many POIs within as certain time span as tourists are. Thus, we might also see differences into their behavioral aspects producing the human photo trails studies in this article. While we have focused on an aggregated view in this focus, the distinction between tourists and residents might be an additional highly interesting layer which we leave open for future work.

# 7. CONCLUSION

In this paper, we investigated and compared a set of hypotheses about urban photo trails across different cities. Towards that end, we employed the Bayesian HypTrails approach to study the sequences of locations generated by users who are uploading photos to the online platform Flickr. For the informed specification of hypotheses, we employed additional data sources such as DBpedia and view counts of Wikipedia articles. Our results suggest that cities share common characteristics, i.e., that the overall ordering of explanations for different cities is quite stable. At the same time, we observe differences on a more detailed level. For example, while proximity is a good explanation across all cities, for Los Angeles we observe movement patterns on a different scale.

In future work, our study can be extended to include other cities. Additionally, it would be interesting to expand the current city-level analysis to a larger scale, e.g., trails that reach across different cities and countries. Furthermore, we aim at a more general framework that allows to capture and integrate different data sources and hypotheses for continuous geo-spatial data into the Markov chain models of the HypTrails approach. Finally, improved tool support for the interactive exploration of location sequences and hypotheses would be a helpful for researchers.

# 8. ACKNOWLEDGMENTS

This work was partially funded by the DFG in the research project "PoSTs II". Special thanks goes to Georg Dietrich from University of Wuerzburg for crawling the Flickr data.

# References

[1] J. An, D. Quercia, and J. Crowcroft. Partisan sharing: facebook evidence and societal consequences. In *Conference on Online Social Networks*, pages 13–24. ACM, 2014.

[2] T. Berners-Lee and M. Fischetti. *Weaving the Web: The original design and ultimate destiny of the World Wide Web by its inventor.* HarperInformation, 2000.

[3] J. Borges and M. Levene. Evaluating variable-length markov chain models for analysis of user web navigation sessions. *IEEE Transactions on Knowledge and Data Engineering*, 19(4):441–452, Apr. 2007.

[4] S. Brin and L. Page. The anatomy of a large-scale hypertextual web search engine. In *International Conference*

*on World Wide Web*, pages 107–117. Elsevier Science Publishers B. V., 1998.

[5] D. P. Brumby and A. Howes. Good enough but i'll just check: Web-page search as attentional refocusing. In *International Conference on Cognitive Modeling*, pages 46–51, 2004.

[6] L. D. Catledge and J. E. Pitkow. Characterizing browsing strategies in the world-wide web. *Computer Networks and ISDN Systems*, 27(6):1065–1073, 1995.

[7] M. Cha, A. Mislove, and K. P. Gummadi. A measurement-driven analysis of information propagation in the flickr social network. In *International Conference on World Wide Web*, pages 721–730. ACM, 2009.

[8] M. Chalmers, K. Rodden, and D. Brodbeck. The order of things: activity-centred information access. *Computer Networks and ISDN Systems*, 30(1):359–367, 1998.

[9] E. H. Chi, P. L. T. Pirolli, K. Chen, and J. Pitkow. Using information scent to model user information needs and actions and the web. In *Conference on Human Factors in Computing Systems*, pages 490–497. ACM, 2001.

[10] E. Cho, S. A. Myers, and J. Leskovec. Friendship and mobility: user movement in location-based social networks. In *International Conference on Knowledge Discovery and Data Mining*, pages 1082–1090. ACM, 2011.

[11] D. J. Crandall, L. Backstrom, D. Huttenlocher, and J. Kleinberg. Mapping the world's photos. In *International Conference on World Wide Web*, pages 761–770. ACM, 2009.

[12] M. De Choudhury, M. Feldman, S. Amer-Yahia, N. Golbandi, R. Lempel, and C. Yu. Automatic construction of travel itineraries using social breadcrumbs. In *Conference on Hypertext and Hypermedia*, pages 35–44. ACM, 2010.

[13] F. Girardin, F. Calabrese, F. D. Fiore, C. Ratti, and J. Blat. Digital footprinting: Uncovering tourists with user-generated content. *Pervasive Computing*, 7(4):36–43, 2008.

[14] F. Girardin, F. D. Fiore, C. Ratti, and J. Blat. Leveraging explicitly disclosed location information to understand tourist dynamics: a case study. *Journal of Location Based Services*, 2(1):41–56, 2008.

[15] M. C. Gonzalez, C. A. Hidalgo, and A.-L. Barabasi. Understanding individual human mobility patterns. *Nature*, 453(7196):779–782, 2008.

[16] B. A. Huberman, P. L. T. Pirolli, J. E. Pitkow, and R. M. Lukose. Strong regularities in world wide web surfing. *Science*, 280(5360):95–97, Mar 1998.

[17] J. Lehmann, R. Isele, M. Jakob, A. Jentzsch, D. Kontokostas, P. N. Mendes, S. Hellmann, M. Morsey, P. van Kleef, S. Auer, et al. Dbpedia–a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic Web*, 2014.

[18] K. Lerman and L. Jones. Social browsing on flickr. In *International Conference on Web and Social Media*, 2006.

[19] F. Mahdisoltani, J. Biega, and F. Suchanek. Yago3: A knowledge base from multilingual wikipedias. In *7th Biennial Conference on Innovative Data Systems Research*. CIDR 2015, 2014.

[20] C. Marlow, M. Naaman, D. Boyd, and M. Davis. Ht06, tagging paper, taxonomy, flickr, academic article, to read. In *Conference on Hypertext and Hypermedia*, pages 31–40. ACM, 2006.

[21] Y. Matsubara, Y. Sakurai, C. Faloutsos, T. Iwata, and M. Yoshikawa. Fast mining and forecasting of complex time-stamped events. In *International Conference on Knowledge Discovery and Data Mining*, pages 271–279. ACM, 2012.

[22] A. Mislove, H. S. Koppula, K. P. Gummadi, P. Druschel, and B. Bhattacharjee. Growth of the flickr social network. In *Workshop on Online Social Networks*, pages 25–30. ACM, 2008.

[23] A. Noulas, B. Shaw, R. Lambiotte, and C. Mascolo. Topological properties and temporal dynamics of place networks in urban environments. *arXiv:1502.07979 [cs.SI]*, 2015.

[24] B. J. Pierce, S. R. Parkinson, and N. Sisson. Effects of semantic similarity, omission probability and number of alternatives in computer menu search. *International Journal of Man-Machine Studies*, 37(5):653–677, 1992.

[25] P. L. T. Pirolli and S. K. Card. Information foraging. *Psychological Review*, 106(4):643–675, 1999.

[26] P. L. T. Pirolli and J. E. Pitkow. Distributions of surfers' paths through the world wide web: Empirical characterizations. *World Wide Web*, 2(1-2):29–45, Jan 1999.

[27] A. Popescu and G. Grefenstette. Deducing trip related information from flickr. In *International Conference on World Wide Web*, pages 1183–1184. ACM, 2009.

[28] A. Popescu, G. Grefenstette, et al. Mining user home location and gender from flickr tags. In *International Conference on Web and Social Media*, 2010.

[29] J. Rekimoto, T. Miyaki, and T. Ishizawa. Lifetag: Wifi-based continuous location logging for life pattern analysis. In *International Conference on Location and Context Awareness*, volume 2007, pages 35–49, 2007.

[30] R. Sen and M. Hansen. Predicting a web user's next access based on log data. *Journal of Computational Graphics and Statistics*, 12(1):143–155, 2003.

[31] B. Sigurbjörnsson and R. Van Zwol. Flickr tag recommendation based on collective knowledge. In *International Conference on World Wide Web*, pages 327–336. ACM, 2008.

[32] P. Singer, D. Helic, A. Hotho, and M. Strohmaier. Hyptrails: A bayesian approach for comparing hypotheses about human trails. In *International Conference on World Wide Web*, 2015.

[33] P. Singer, D. Helic, B. Taraghi, and M. Strohmaier. Detecting memory and structure in human navigation patterns using markov chain models of varying order. *PloS one*, 9(7):e102070, 2014.

[34] P. Singer, T. Niebler, M. Strohmaier, and A. Hotho. Computing semantic relatedness from human navigational paths: A case study on wikipedia. *International Journal on Semantic Web and Information Systems*, 9(4):41–70, 2013.

[35] R. W. Sinnott. Virtues of the haversine. *Sky and Telescope*, 68(2):158, 1984.

[36] C. Song, Z. Qu, N. Blumm, and A.-L. Barabási. Limits of predictability in human mobility. *Science*, 327(5968):1018–1021, 2010.

[37] C. H. Tai, D.-N. Yang, L. T. Lin, and M. S. Chen. Recommending personalized scenic itinerarywith geo-tagged photos. In *International Conference on Multimedia and Expo*, pages 1209–1212. IEEE, 2008.

[38] S. Walk, P. Singer, and M. Strohmaier. Sequential action patterns in collaborative ontology-engineering projects: A case-study in the biomedical domain. In *International Conference on Conference on Information & Knowledge Management*. ACM, 2014.

[39] R. West and J. Leskovec. Human wayfinding in information networks. In *International Conference on World Wide Web*, pages 619–628. ACM, 2012.

[40] R. West, J. Pineau, and D. Precup. Wikispeedia: An online game for inferring semantic distances between concepts. In *International Joint Conference on Artifical Intelligence*, pages 1598–1603. Morgan Kaufmann Publishers Inc., 2009.

[41] R. W. White and J. Huang. Assessing the scenic route: measuring the value of search trails in web logs. In *Conference on Research and Development in Information Retrieval*, pages 587–594. ACM, 2010.

[42] S. E. Wiehe, A. E. Carroll, G. C. Liu, K. L. Haberkorn, S. C. Hoch, J. S. Wilson, and J. D. Fortenberry. Using gps-enabled cell phones to track the travel patterns of adolescents. *International Journal of Health Geographics*, 7(1):22, 2008.

[43] J. Yang, J. McAuley, J. Leskovec, P. LePendu, and N. Shah. Finding progression stages in time-evolving event sequences. In *International Conference on World Wide Web*, pages 783–794. ACM, 2014.