Noname manuscript No.
(will be inserted by the editor)

Exploiting social media information towards a context-aware recommendation system

Michalis Korakakis · Evaggelos Spyrou · Phivos Mylonas · Stavros J. Perantonis

Received: date / Accepted: date

Abstract The rise of the social networks during the last few years has provided a vast amount of knowledge in several domains. Among them, route planning and point-of-interest recommendation have significantly benefited. Seen from the side of a tourist, they consist two challenging and time-consuming tasks since they may rely on many parameters and are limited by several constraints, such as time and budget available, user preferences etc. In this paper we present Xenia, a context-aware system that works towards solving the aforementioned problems. More specifically, it aims to automatically construct travel routes, i.e., ordered visits to various points-of-interest. The user (tourist) indicates an initial and an ending point and her/his available time budget and the system proposes travel routes that maximize her/his travel experience, while adhering to the aforementioned limitations. This particular route planning problem is widely known as the Tourist Trip Design Problem, having several variations. In this work we solve this problem by modelling it through the Orienteering Problem. We harvest geo-tagged photos from the well-known social network Flickr and using the user-generated textual meta-

M. Korakakis
Department of Informatics, Ionian University, Corfu, Greece
E-mail: michalis.korakakis@gmail.com

E. Spyrou
Institute of Informatics and Telecommunications, National Center for Scientific Research - "Demokritos," Athens, Greece
E-mail: espyrou@iit.demokritos.gr
Tel.: +30-210-650-3175

Ph. Mylonas
Department of Informatics, Ionian University, Corfu, Greece
E-mail: fmylonas@ionio.gr

S.J. Perantonis
Institute of Informatics and Telecommunications, National Center for Scientific Research - "Demokritos," Athens, Greece
E-mail: sper@iit.demokritos.gr

data that accompany them we extract areas-of-interest within a given city along with their semantics. Moreover, by utilizing the timestamps of the photos we are able to identify the trajectory patterns of tourists, to detect popular places-of-interest and finally to estimate the average visit duration. Using this historical data we propose travel routes for 4 of the most popular Greek cities. The effectiveness of our approach is validated upon a twofold validation consisting by a) a comparison vs. the most typical baselines that have adopted by state-of-the-art works and b) an empirical evaluation by real-life users.

Keywords Travel route recommendation · POI extraction · socially-generated knowledge · Flickr

1 Introduction

Back in the era of analog photography, the main reason for which people took photos was to promote communication of information about themselves to themselves and also to others (e.g., relatives, social circle, even future generations) [5]. This information was often related to a single moment, an event (e.g., a wedding), or vacations. To facilitate organization and narration based on the photo content, people often added some hand-written notes behind photos. These often comprised of a place, a date, people present and a brief description of the situation depicted. We may argue that this was indeed a "primitive" version of *tagging*. Photos were also organized into albums, frequently based on the occasion they had been taken. Of course, due to costs needed for supplies and printing, the number of photos taken yearly by the average person was relatively lower compared to the one of the digital era [45].

On the other hand, during the last few years digital photography has definitely dethroned the analog one. Each person is dynamically a photographer, since the majority of modern smartphones is equipped with high quality cameras. This, in addition to the practically zero cost of digital photos has led to a huge increase in the amount of photos taken everyday [9]. Moreover, the continuous expansion of the coverage and speed of broadband mobile networks accompanied by advances in web-based technologies, have dramatically changed the fundamental norms of social interaction. Nowadays, the average person is willing to share her/his photos within social networks and additionally to provide information about her/his "whereabouts" (i.e., the precise or vague location she/he has been) [10, 15]. Typically referred to as "check-in," this information is often linked to some multimedia content (photos and/or videos) and loosely annotated by a set of descriptive keywords (tags).

Adding such "metadata" to multimedia content plays a crucial role in today's research activities, since huge, yet weakly annotated datasets are now offered. More specifically, although users do not tag their digital content having in mind the needs of the research community, the provided information is often characterized as "social knowledge," while users may act as "social sensors." This indicates that some kind of knowledge regarding a specific domain/application may indeed be extracted based on the content provided and annotated by them. During the last decade, significant research efforts have been turned towards Flickr [41] and many research areas such as multimedia content retrieval, tag recommendation, content localization, travel applications, human activity tracking etc. have benefited substantially. Within the context of this work, we deal with the analysis problem of the online "footsteps" that correspond to the actual presence of users in certain places, which is often referred to as "digital footprinting" [14, 1]. These are often processed in order to extract some knowledge about the users' whereabouts, interests or produce personalized recommendations based on their needs.

More specifically, in this work we propose *Xenia*, which is a novel, context-aware system, aiming to automatically construct and propose travel routes to potential tourists. We define a travel route as "an ordered set of POIs, built upon various constraints and parameters." Xenia generates these routes based on socially-generated knowledge derived from the metadata that accompany geo-tagged photos collected from the well-known Flickr[1] social network. We should clarify herein that within this work the visual information of photos is discarded, thus we work only with the available textual, temporal and geospatial metadata of the collected photos.

Upon a clustering procedure on a collection of geo-tagged photos, we are able to discover "Areas-of-Interest" (AOIs), without any prior knowledge of the given urban area. According to Hu et al. [19], an AOI is an "area within an urban environment which attracts people's attention." Based on this definition, and given that the application domain of the problem at hand is tourism, we may argue that an AOI is considered an area that contains tourist attractions such as landmarks, museums, art galleries, places of worship etc. However, as it has been argued [29] for other problems, e.g., when the focus is in local residents, AOIs may contain commercial places, parks, i.e., places where she/he could spend her/his spare time. Of course, the notion of an AOI may differ based on the users' age, gender, social status, culture, nationality, etc. Due to the aforementioned, it should be clear that an AOI may not be strictly defined. Moreover, AOIs are often vague areas with uncertain boundaries while they cannot be referred to by vernacular names due to the co-existence of multiple POIs within them [17]. Furthermore, and to the best of our knowledge, there do not exist any "official" and/or "accurate" lists of AOIs, since it is obvious that accuracy may not be estimated. However, a few such socially-generated lists are publicly available.

In the context of the presented work, an AOI is constructed under the assumption that is an area attracting a large number of visitors (thus "containing" a large number of geo-tagged photos). Each AOI should contain one or more "Places-of-Interest" (POIs), since a single POI is often not adequate to characterize an AOI. We should clarify that a POI may be a *single* attraction either limited to a relatively small geographic area (e.g., a statue, or a small building) or to relatively large one (e.g., a museum or a monument). We adopt a clustering approach in order to discover AOIs as distinct high density areas

[1] `http://www.flickr.com`

extracted from raw sets of geo-tagged photos collected from Flickr, thereby eliminating any noise by treating it as low density outlier areas, i.e., those that do not correspond to AOIs/POIs. Upon selecting an AOI, we then try to identify any potential POIs that it may contain, by ranking textual metadata of photos, cross-checked to public geo-information related databases, e.g., map services. Given a set of extracted POIs we then provide an empirical metric, so as to quantify the "gain" a user has by visiting them.

Since our goal is to identify and utilize trajectory patterns derived from tourists for the purpose of measuring the popularity that a POI has along with its average visit duration, we feel that Flickr is indeed an appropriate choice, as a) it provides a powerful API;[2] b) the majority of its hosted photos and their accompanying metadata may be used for non-profit activities;[3] c) photos hosted within Flickr are usually *geo-tagged*, i.e., the location of the depicted content has been added either automatically (e.g., by the camera/smartphone used) or manually (e.g., by the photographer); and d) the majority of cameras used by Flickr users are smartphones and consumer or entry level SLR cameras.[4]

The proposed system tries to satisfy the following user requirement: "Given a starting and an ending location and the amount of time available, specify a travel route that maximizes the user's gain, based on a set of predefined criteria." We model this particular problem as a Tourist Trip Design Problem (TTDP) [40] and solve it through a variant of the Orienteering Problem (OP) [47]. More specifically, by solving the corresponding OP, we are able to create travel routes that may adhere to a multitude of different constraints and parameters, imposed either explicitly by the user or indirectly through various trip-related limitations. We should emphasize that the presented approach does not rely on any manually constructed ground truth (e.g., of POIs within the given examined urban area) and contrary to other works is fully-automated.

The remaining of the paper is organized as follows: Section 2 presents related research works in the fields of general recommendations using socially-generated knowledge, extracted from Flickr metadata and of recommendation problems that have been tackled as variations of the TTDP. Then, in section 3 we present the OP along with its integer programming formulation that we subsequently exploit for the purpose of solving the herein presented variation of the TTDP. The proposed system, i.e., Xenia is then presented in section 4, where we discuss in detail the set of algorithms involved, i.e., geo-tagged photo clustering and AOI extraction, tag ranking, POI detection and Travel Substreams identification. Extensive experimental results are presented in section 5. Finally, conclusions are drawn in section 6, where we also discuss possible further extensions of Xenia and plans for related future work.

[2] `https://www.flickr.com/services/api/`

[3] `https://www.flickr.com/creativecommons/`

[4] `http://blog.flickr.net/en/2015/01/13/camera-ownership-on-flickr-2013-2014/`

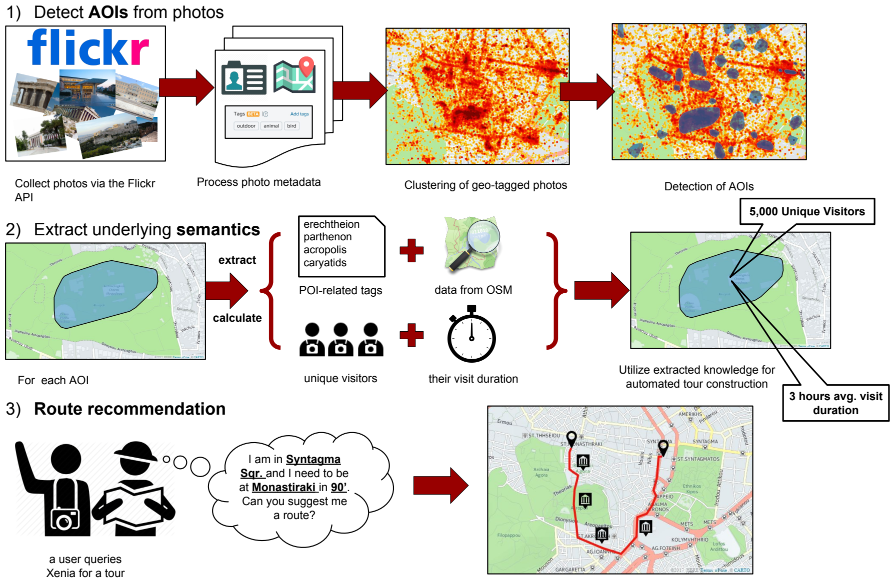

Fig. 1 A visual overview of the proposed approach.

2 Related Work

As it has already been discussed in section 1, Flickr has been facilitating the needs of the research community by making available a huge amount of geo-tagged and weakly annotated photos. Thus, it has been widely used for research purposes in many areas such as tourism and/or personalized POI recommendations. In this section we provide related work in the aforementioned areas that uses (meta)data collected from Flickr. Moreover, since the main focus of this paper is to provide a solution to the TTDP, related work is extensively presented.

2.1 Tourism Recommendations using Social Knowledge from Flickr

Typical applications in this area focus on automatically discovering main attractions, letting users decide which to visit and/or also aim to organize the users' schedule and help them visit as many attractions as they wish in a time efficient way. Van Canneyt et al. [44] tackled the recommendation problem by discovering trends and ranking POIs based on their popularity and other temporal information. POIs were discovered upon spatial clustering and ranked based on their popularity, the context of their users and temporal information. The user selected a few places of interest and timeslots she/he was available and their system proposed the best coverage of these. Similarly, Kisilevich et al. [23] also used spatial clustering to discover POIs and then, for each resulting cluster they described its temporal dynamics upon a time-series analysis.

They characterized clusters in a spatio-temporal way as stationary, reappearing, occasional and regular moving. Cao et al. [4] also included visual features of images in the process, by extracting a set of representative images and tags for each cluster. They proposed a tourism recommendation system which responded to users' photo queries with suggestions. Popescu and Grefenstette [38] exploited the temporal aspect of metadata so as to make estimations regarding visiting times for attractions and deduced what a tourist is able to visit in a city within a day. Jain et al. [20] extracted routes that start from a given location and proposed one that visits popular places using certain distance constraints, however without considering time needed. Popescu et al. [39] provided an extension and generated new trips by combining those mined. Hao et al. [16] recommended popular places of a specific region, annotate some aspects of them (e.g., as landmarks or activities) and summarized landmarks by providing representative images. Sun et al. [42] also identified landmarks which then were ranked using their popularity and minimum distances with maximum popularity. Finally, Jiang et al. [21] extracted user preference topics (e.g., cultural, cityscape, or landmark) based on textual metadata so as to find similar users and recommended locations based on them.

2.2 Tackling the TTDP problem with Social Knowledge

The field of trip recommendations has become quite popular during the last years with numerous works that typically aim to solve variations of the TTDP using socially-generated knowledge and focusing on user metadata and/or whereabouts, shared in popular social networks. One of the early works towards the aforementioned direction was the one of De Choudhury et al. [8] where the problem was modelled as a single graph of POIs, their popularities and the transit times among them. Another early systematic approach was also the one of Lu et al. [30], who presented a system, namely *Photo2Trip* which recommended routes within a city and also "internal" routes, i.e., bounded within a smaller area. Typical constraints in the process of selecting a valid route have been combined by Lu et al. [31], which considered the available time, budget and number of POIs. Moreover, Wu et al. [48] incorporated knowledge of weather predictions to dynamically adapt routes, accordingly. Also, Vansteenwegen et al. [46] took into account opening hours of attractions and created personalized travel routes according to users explicit preferences that span among several days. A lot of effort has been devoted on personalization issues. Majid et al. [33] tried to "transfer" knowledge, by using user preferences from a given city to recommend locations to another city. Kurashima et al. [24] proposed an approach that models the photographer's behavior and upon a combination with popular routes between landmarks it was able to recommend personalized routes. Quercia et al. [35] modeled the human perception of locations as beautiful, quiet, and happy and proposed routes that balanced between short and pleasant. Chen et al. [7] considered specific user profiles (e.g., gender, age etc.) and also group types (e.g., couple, family, etc.)

in the process. At a next step, Lim et al. [26] clustered tourists into tour groups in a way that each would contain tourists with similar interests. Then they recommended routes per group, matching its interests, while they also assigned an appropriate group guide. Additional knowledge was incorporated in the process by Brilhante et al. [2], who mined tourist routes from Flickr and matched them on POIs harvested by Wikipedia and proposed personalized tourist routes that may span into multiple days. Yoon et al. [50] did not rely on users' digital whereabouts and instead used trajectories that had been generated by volunteers carrying GPS logger devices. Using these trajectories they detected "stay-points," which were then used to select POIs and propose routes. Hsieh et al. [18] did not rely on moving trajectories but solely on timestamped route data, at an effort to determine the best time to visit a given place. The TTDP is often modelled as an OP, e.g., as in the work of Lim et al. [25]. Other heuristics that may be involved within this process are the explicit selection of a subset of POIs [49] or ordering constraints [13] set by the user and also interaction between users and user history [12]. Research is currently heading to providing full tourist packages that may combine stay, transportation and sightseeing e.g., as the work of Lu et al. [32]who extracted possible means of transportation among POIs. We feel that the proposed work is close to the one of [25], however we introduce extracted semantics of AOIs in the process.

3 The Orienteering Problem

The Orienteering Problem (OP) is based on the orienteering sport game, in which several locations with an associated score have to be visited within a given time limit [6]. It is also known as the Selective Traveling Salesman Problem (STSP), the Maximum Collection Problem (MCP) and the Bank Robber (BR) problem. Also, it is related to the Knapsack combinatorial optimization problem [37]. The classical OP [47] is typically defined as: "Given a set of vertices, each assigned with a score, determine a path bounded in terms of length, maximizing the sum of scores of visited vertices." We should note that this score is typically heuristically determined, depending to the context of the application.

In order to point out its relation to the well-known traditional Traveling Salesman Problem (TSP), the latter may be stated as: "Given a list of cities, an origin city and the distances between each pair, select a route that minimizes the distance travelled and returns to the origin, while each city is visited exactly once." On the contrary and within the same domain, we could state the OP as "Given a list of cities, an origin and a terminal city and the distances between each pair, select a route that maximizes the total gain, earned upon a visit to a city, while also adhering to a positive time budget while each city is visited at most once." Contrary to the TSP, and considering cities as vertices of a graph, interconnected by the edges of the graph, we may observe that within the OP a) not all nodes need to be visited; b) the origin and terminal nodes need not be the same; and c) the criterion to be optimized is different.

Within this work we choose to model the TTDP based on the classical formulation of the computationally NP-Hard OP of [47] and adopt a solution as an integer problem. The goal of OP, adopted to the problem at hand is to find a travel route within an area that, given a starting and a terminal POIs (selected by the user), it maximizes the total score, which is calculated based on a set of heuristics and earned upon when visiting each POI, while also adhering to a positive time budget. To formulate it, let us first assume the existence of a set of vertices $\mathcal{V} = \{v_i, i = 1, 2, \ldots, N\}$. Each vertex v_i corresponds to a POI and has been assigned a non-negative score s_i. We define v_1 and v_N to be the starting and ending vertices, respectively. We also assume that it may not be possible to visit all vertices, by imposing a time constraint $T_{\max}$, and setting a time cost t_{ij} between vertices v_i and v_j. The goal is to visit a subset of $\mathcal{V}$, so as a) to maximize the sum of scores of visited vertices; and b) visit each vertex at most once.

Using the aforementioned notation, we may formulate the OP as an integer problem. Variable x_{ij} is equal to 1, when a visit to v_i has been followed by a visit to v_j, else is equal to 0. Also, u_i provides the position of v_i within the given path. The objective function (score) that has to be maximized is defined by Eq. 1.

$$\max \sum_{i=2}^{N-1} \sum_{j=2}^{N} s_i x_{ij} \tag{1}$$

given the following set of six constraints (Eq. 2-7), namely:

$$\sum_{j=2}^{N} x_{1j} = \sum_{i=1}^{N-1} x_{iN} = 1 \ , \tag{2}$$

which guarantees that the extracted path's starting and ending points are indeed v_1 and v_N, respectively, since e.g., if v_1 is the starting point, exactly one vertex $v_j,\ j = 2, \ldots, N$ will be visited (assuming that a route may contain more than one vertices) and similarly, if v_N is the ending point, it will be visited from exactly one vertex $v_j,\ j = 1, \ldots, N-1$. In combination with the following Eq. 4, it is guaranteed that v_1 may not be visited by another vertex, while starting from v_N, another vertex may not be visited.

$$\sum_{i=1}^{N-1} x_{ik} = \sum_{j=2}^{N} x_{kj} \leq 1; \forall k = 2, \ldots, N-1 \ , \tag{3}$$

which guarantees the connectivity of the path, i.e., the result is a tree and not a forest. It also ensures that each vertex is visited at most once, and from each vertex, at most one other vertex may be visited.

$$\sum_{i=1}^{N-1} \sum_{j=2}^{N} t_{ij} \leq T_{\max} \ , \tag{4}$$

which obviously imposes the time constraint T_{max} on the travel route, by summation of the individual time costs t_{ij} for the visit from vertex v_i to v_j and for all visited vertices.

$$2 \leq u_i \leq N; \forall i = 2, \ldots, N \tag{5}$$

and

$$u_i - u_j + 1 \leq (N-1)(1 - x_{ij}); \forall i, j = 2, \ldots, N \ , \tag{6}$$

have been proposed by Miller et al. [34] and are used to eliminate subtours, i.e., directed cycles within the constructed route, or routes with larger length than N. Finally,

$$x_{ij} \in \{0, 1\}; \forall i, j = 1, \ldots N \ , \tag{7}$$

i.e., as it has already been mentioned, x_{ij} is equal to either 0, or 1, depending on the existence of a directed edge from v_i to v_j in the final path. Additionally and as expected, we assume that $t_{ij} = t_{ji}$, i.e., equal travel time among the corresponding vertices.

4 Xenia

In this section we shall present our context-aware trip recommender system, namely "Xenia." The proposed system follows the well-known travel route recommendation paradigm, i.e., it constructs travel substreams based on a set of photos that a user has taken and utilizes the extracted knowledge for the purpose of tailoring a multitude of trip related parameters for a specific urban region in a similar manner to that of [8, 25, 2]. It differentiates since it is able to exploit socially-generated knowledge, i.e., user-generated, geo-tagged metadata from Flickr, in order to discover AOIs and subsequently identify the POIs that they may contain.

4.1 AOI Extraction

The goal of the first step of Xenia towards route recommendation is to discover AOIs within an urban area "in the wild", i.e., without any prior knowledge. To this goal it uses geospatial information extracted from photos that have been collected from Flickr. Within the context of this work, we shall assume that an AOI is characterized both by a large number of attracted visitors and by a relatively large number of POIs, contained within. Since the geospatial information of a photo is expressed as a 2D point (i.e., the corresponding latitude and longitude of the location that it has been taken[5]), the task of AOI identification may be tackled by relying on the notion of density. Density-based clustering techniques are able to detect distinct clusters that are defined

[5] To be more accurate, this geospatial information, when it is manually generated by the users, is prone to errors, since geo-tagging may in some cases be a subjective task. Thus in some cases it represents the location where a photo has been *tagged*.

as dense regions having an arbitrary size and shape and are surrounded by regions of low-density. This particular trait can also be encountered in the case of urban AOIs [19].

Moreover, density-based methods can detect automatically the either optimal or desired or appropriate number of clusters, depending on the application's context, without relying on explicit user input. The aforementioned characteristic is rather important since the number of AOIs that can be spotted within a region is based on a multitude of different factors that could widely vary among different geographical regions (e.g., depend on their cultural, economical and geospatial characteristics) and thus its accurate prediction is a complex and difficult task for a human to perform. Finally, this particular family of clustering methods is robust to noise. Within this work, noise is perceived as secluded ("outlier") areas containing a rather small amount of geo-tagged photos.

Therefore, we choose to cluster these data using HDBSCAN [3], an extension of the well-known density-based DBSCAN algorithm [11]. DBSCAN's main weakness is its inability to detect clusters with varying densities, a disadvantage that is effectively solved by HDBSCAN. We should emphasize that this attribute is crucial for this work, since it is rather usual for a particular urban region to have a limited set of AOIs that attract huge numbers of crowds due to the POIs that they contain (e.g., *Parthenon*, *Temple of Olympian Zeus* etc., in the case of Athens), whereas the rest AOIs are less popular. In order to achieve this, HDBSCAN converts DBSCAN to a hierarchical clustering algorithm, so as to extract a set of significant density-based clusters. More specifically, by using the generated dendrogram (i.e., a structural graphical representation of the distances between the connected components of the clustered data) it becomes possible to obtain clusters with different density thresholds (i.e., in a similar manner to multiple executions of DBSCAN, each with a diverse set of parameters). The process of obtaining a flat clustering is accomplished by introducing the concept of "cluster stability" (i.e., a measurement of the persistence of a cluster at various levels of the constructed hierarchy) and thus treating the task of cluster extraction as an optimization problem, which objective is to find those that maximize the overall value of the aforementioned metric.

We should emphasize that HDBSCAN requires a single user-specified parameter as its input: the minimum size of points (hereafter referred to as `minPts`) that a given cluster may contain. In this work, we implement HDBSCAN using ELKI [36], an open-source data mining software framework.[6] We follow a trial-and-error approach and experiment with various parameter values, using relatively high and low values for `minPts`. Our goal is to identify clusters that correspond to a potential "popular" area (i.e., an area visited by a large number of tourists) while at the same time they do not cover a large

[6] In particular, we use the HDBSCAN-SLINK version, which differs from the original algorithm due to using SLINK instead of Prim's algorithm for the purpose of obtaining a single-linkage dendrogram.

part of the examined urban region, since this way they may contain a very large number of POIs, which is an undesired property.

4.2 Tag Ranking

At the next step, the goal is to find a set of semantically meaningful tags that can uniquely characterize the POIs contained within a particular geo-cluster, by adopting the following tag-ranking approach: We identify the top k-representative tags using the $N \times N$ co-occurrence matrix C for each cluster. $C(i, j)$ indicates the number of photos where the i-th word co-occurs with the j-th word, i.e., they have been both used to tag it. We should note that N is the number of the unique words (tags) that are contained within the photos of the examined cluster. Through the use of this matrix, we are capable to pair together tags that are semantically related (e.g., *acropolis* and *parthenon*). By employing this technique we are able to accumulate a varied set of tags that can be exploited so as to distinguish a POI, while ignoring those with high frequencies.

4.3 POI Extraction

Upon selecting a ranked set of tags, the next step attempts to extract the various POIs that a given geo-cluster may contain. To achieve this we collect relevant information (i.e., names, categories and geographic coordinates) about the POIs of the examined urban area from OpenStreetMap,[7] a free-to-use worldwide map that relies on crowd-sourced volunteered geographic information. In order to discern whether a specific POI belongs to a geo-cluster we choose to use their Levenshtein distance [27], between the set of words that characterize it and the top-k representative tags for the geo-cluster. In principle, the Levenshtein distance is a well-known, widely used string metric for measuring the difference between two sequences. Between two given words, it is defined as the minimum number of edits needed to transform one word into the other, with the allowable edit operations being insertion, deletion or substitution of a single character. The result of this process is the identification of a set of POIs within each cluster.

4.4 Travel Substreams

Since the accompanying metadata of a photo that has been collected from Flickr includes a distinctive name of the user who has uploaded it, along with the exact date and time the photo had been taken,[8] we are able to divide user

[7] `https://www.openstreetmap.org`

[8] Under the assumption that the photographer has correctly set the date on her/his camera or phone.

travel histories into "substreams." In other words, for each user we construct her/his travel substreams, i.e., consequential visits to POIs. This is the first step for the computation of the average visit duration for a specific POI. We define a travel route r as a sequence $\{p_1, p_2, ...p_N\}$ of photos. All route members should comply to the following rules:

- $p_1.\texttt{date} \leq p_2.\texttt{date} \leq \ldots \leq p_N.\texttt{date}$, i.e., all members should be consecutive;
- $p_1.\texttt{user} = p_2.\texttt{user} = \ldots = p_N.\texttt{user}$, i.e., all members should belong to the same user;
- let p_i, p_j denote a pair of consecutive route members. Then $p_i.\texttt{date} < p_j.\texttt{date} + T_h$, i.e., a restriction is imposed on the maximum time difference between photos in order to be considered as consecutive within a route.

At this point we should note that without loss of generality we associate a geo-tagged photo to a certain POI by measuring their proximity, hence we only consider photos that are taken within a radius of 100m from the geo-location of the POI. Using the whole dataset, and for each unique user we are able to construct a set of her/his travel substreams, according to the aforementioned rules. The visit duration per user for a given POI is calculated as the temporal difference of the first and the last photos she/he had taken within it.

4.5 Scoring Functions

The score function used for the selection of a POI is heuristically determined using the corresponding distinct number of visitors for a POI, i.e., its *popularity*. We consider that a high-traffic POI is of value for a tourist, based on the "wisdom-of-the-crowd", i.e., the general opinion that a group of people possesses. Moreover, for each POI we consider the *average visit duration*, calculated based on the visit durations of all users that have visited this specific POI. We solve the integer formulation of the OP using the Gurobi Optimizer,[9] a solver for mathematical programming problems. Finally, in order to determine real-life pairwise travel times between POIs we use the Google Maps Distance Matrix API.[10]

5 Experiments

In this section we present all performed experiments regarding the AOI, POI extraction, the travel substreams construction and the evaluation of the proposed system, using a dataset collected from Flickr.

[9] `http://www.gurobi.com/products/gurobi-optimizer`
[10] `https://developers.google.com/maps/documentation/distance-matrix/`

Table 1 Locations, areas and statistics from the collected sets of photos, for the 4 Greek cities used throughout the experimental evaluation of the proposed system.

City	South West	North East	Area (km^2)	# of photos	# of unique tags	# of users	# of routes	avg.# of POIs /user
Athens	37.9488, 23.6869	38.0334, 23.7898	84.88	150645	31078	6519	9468	342.9
Thessaloniki	40.5862, 22.8995	40.6526, 22.9890	55.80	28238	11723	1365	1949	64.9
Heraklion	35.1484, 24.9453	35.3521, 25.2053	535.45	17119	5444	1130	1195	91.7
Chania	35.3297, 23.9098	35.6018, 24.2016	800.98	26013	6443	1480	1636	106.4

5.1 Data Sets

In order to provide both a qualitative and quantitative evaluation of Xenia, we selected 4 popular Greek cities of high touristic interest and various sizes. More specifically, we applied the proposed system on Athens,[11] Thessaloniki,[12] Heraklion,[13] and Chania.[14] In the case of Athens and Thessaloniki, we made the valid assumption that a potential tourist will not use a car, but rather she/he will prefer to move on foot, due to the large pedestrian touristic areas available. Considering that a vast amount of popular POIs from these aforementioned regions originate in the city centers and the surrounding areas, we executed a geo-query so as to collect photos solely from these municipalities and their nearby suburbs and exclude popular yet mundane districts for tourists (e.g., Athens airport, Piraeus port, etc.) that are of no value to them. In order to define the bounding boxes which were used for the necessary queries to the Flickr API, we used information gathered from the Geodata.gov.gr[15] website, which serves as the national open data catalogue for Greece. As a result of this process, we ended-up with 150645 and 28238 photos, respectively. On the other hand and in case of Heraklion and Chania, we assumed that tourists typically rent a car, therefore are able to visit more distant POIs. Corresponding geo-queries covered the municipal areas of these two cities, ending up with 17119 and 26013 photos, respectively. In Figs. 2a–5a we illustrate the aforementioned areas that are summarized in Table 1.

All collected photos are geo-tagged, dated between January 2004–December 2015 and acquired from Flickr social network using its public API.[16] Then through the use of a manually created stoplist and regular expressions, we removed tags that either did not have a semantic relation to the respective photo that they were attached to (i.e., automatically added tags by smartphones and cameras e.g., *iphone*, *android*, etc.) or were too generic (i.e., tags that are both common and spread to the whole area, thus not providing any useful information, while also tending to be amongst the most popular e.g., *holidays*, *Greece*, *Athens*, etc.). Upon completion of this process and e.g. in

[11] `https://en.wikipedia.org/wiki/Athens`
[12] `https://en.wikipedia.org/wiki/Thessaloniki`
[13] `https://en.wikipedia.org/wiki/Heraklion`
[14] `https://en.wikipedia.org/wiki/Chania`
[15] `http://geodata.gov.gr`
[16] `https://www.flickr.com/services/api/`

case of Athens, given the initial 40487 tags, we ended up with overall 31078 unique ones. The number of unique tags per city are also summarized in Table 1. Moreover, in Figs. 2b–5b we present density-based visualizations of the distributions of the collected photos in the whole areas used throughout our experiments and for each city. Intuitively, having empirical knowledge e.g., upon a visit to these cities, one should notice that high-density areas correspond to places of increased touristic importance.

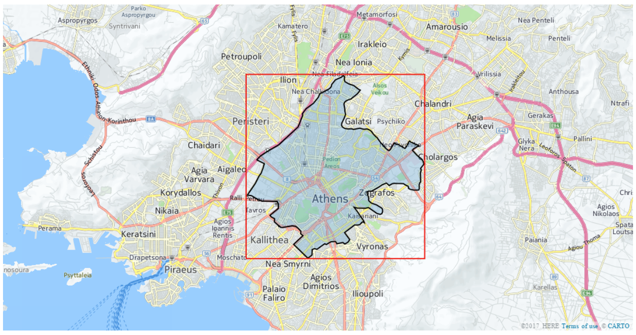

(a) Selected municipal region.

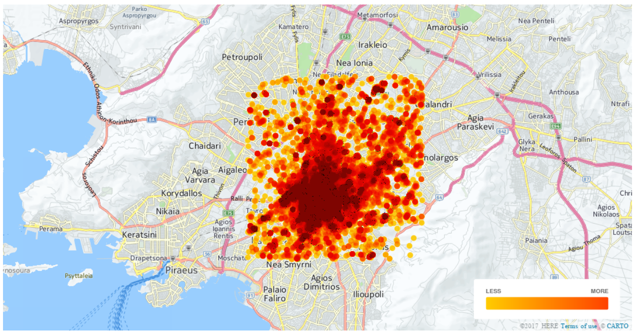

(b) Density-based visualization of Flickr photos.

Fig. 2 The selected municipal region of Athens and the density-based visualizations of Flickr photos that lie within the corresponding geographic bounding box

5.2 AOI, POI extraction and travel substreams construction

Due to the variety of the number of photos taken among the selected cities and also since the areas of Chania and Heraklion are 10× the areas of Athens and Thessaloniki, respectively, finding a common value for `minPts` appears a difficult task and also is not compulsive. Thus, and upon the process that has been described in subsection 4.1, we set `minPts` = 250 for Athens and Thessaloniki and 100 for Chania and Heraklion. Results upon the application of HDBSCAN on each city are depicted in Table 2.

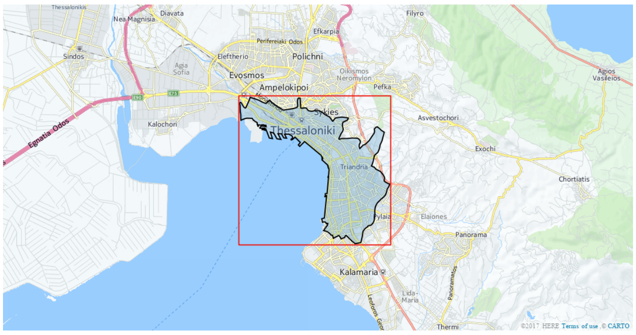

(a) Selected municipal region.

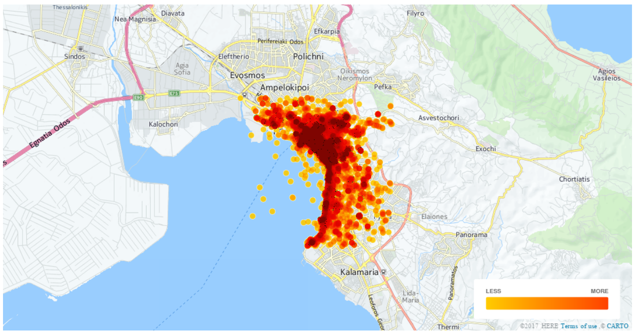

(b) Density-based visualization of Flickr photos.

Fig. 3 The selected municipal region of Thessaloniki and the density-based visualizations of Flickr photos that lie within the corresponding geographic bounding box

Table 2 The choice of `minPts` and the number of extracted AOIs and POIs per city.

city	`minPts`	**AOIs**	**POIs**
Athens	250	69	42
Thessaloniki	250	20	24
Heraklion	100	32	15
Chania	100	51	18

The process of POI recognition and extraction, presented in subsection 4.2, lead to the identification of a number of places. One may notice that the number of places is smaller than the one of the AOIs (geo-clusters). This happens since we filtered all places that do not correspond to landmarks (e.g., restaurants, bars, cinemas, etc.), yet are frequently visited by tourists. Moreover, it is possible for an AOI to simply act as an area that provides a scenic view due to its advantageous location without actually containing any POIs (e.g., Mtn. Lycabettus in Athens). We first queried the dataset for each user separately. Then we ordered all photos per user at ascending date. As we mentioned in subsection 4.4, we set $T_h = 6h$, under the assumption that consecutive photos taken at very large intervals, belong to different sequences. We ended up

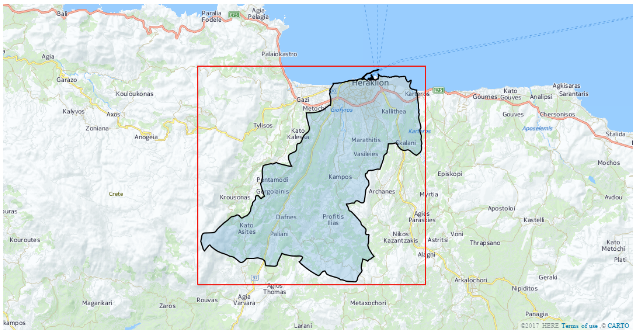

(a) Selected municipal region.

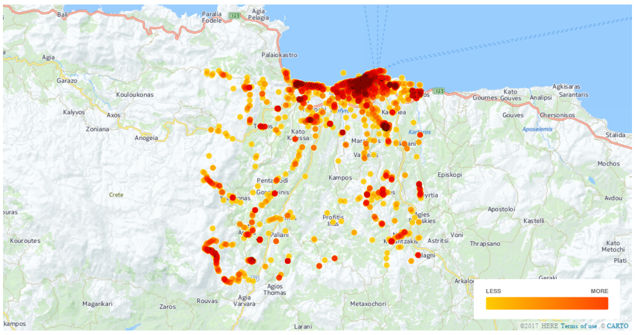

(b) Density-based visualization of Flickr photos.

Fig. 4 The selected municipal region of Heraklion and the density-based visualizations of Flickr photos that lie within the corresponding geographic bounding box

with this conclusion upon careful inspection of the available data set. We then discarded sequences that consisted of less than 3 members.

In Fig. 6 we illustrate the extracted AOIs for all cities, overlaid to the densities of the photos (which have been presented in Figs. 2b–5b. Empirically, one may intuitively verify the clustering process based on her/his knowledge on the aforementioned cities.

5.3 Evaluation

For the evaluation of the proposed system, we opted for a twofold approach: a) an objective evaluation by using a set of well-known baselines from the state-of-the-art and appropriate information retrieval metrics and b) a user-oriented subjective evaluation, i.e., aiming to investigate whether real-life users are satisfied by the proposed travel routes.

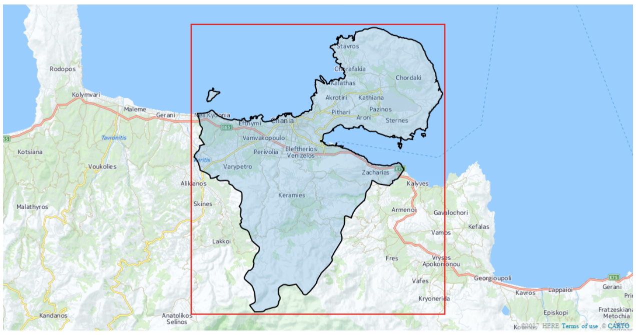

(a) Selected municipal region.

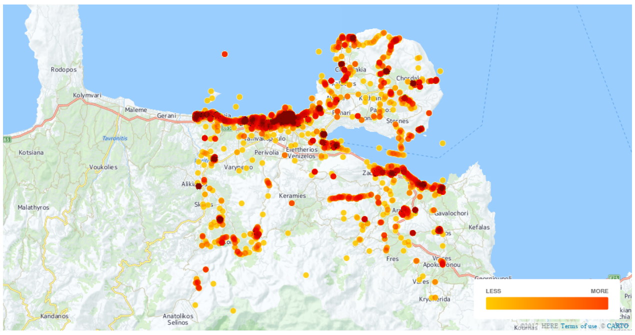

(b) Density-based visualization of Flickr photos.

Fig. 5 The selected municipal region of Chania and the density-based visualizations of Flickr photos that lie within the corresponding geographic bounding box

5.3.1 Quantitative Evaluation

The main objective of the proposed system is to recommend travel routes that can actively contribute to the increase of a tourist's overall travel experience for a holiday destination that is characterized by a large number of POIs. Thus, we assess the performance of our proposed system against the following baselines, which are in general aligned to the ones used in the state-of-the-art [25, 2]:

– *Random POI Selection (RPOI).* RPOI randomly selects an unvisited POI at each step and adds it in the route.
– *Greedy POI Selection (GPOI).* GPOI chooses the next POI according to its corresponding earned score, i.e., the highest score from the set of unvisited POIs.
– *Nearest POI Selection (NPOI).* NPOI constructs the travel route by adding an unvisited POI to the latter, according to its distance from the POI that was selected in the immediate preceding iteration of the process.

We chose to evaluate Xenia using the well-known Precision and Recall metrics, which in our case are defined and applied as:

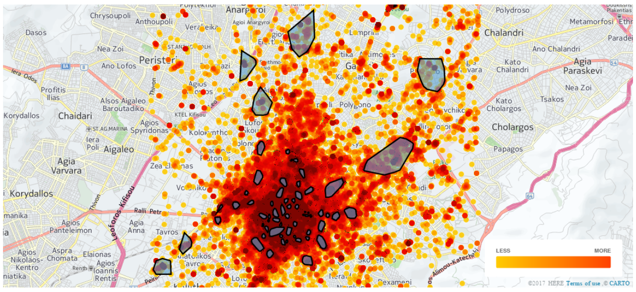

(a) Athens

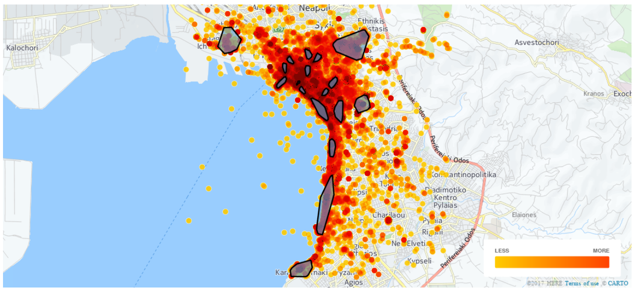

(b) Thessaloniki

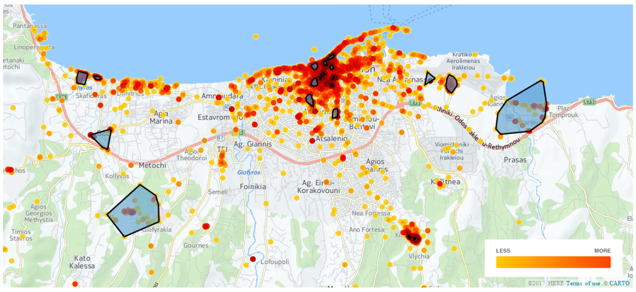

(c) Heraklion

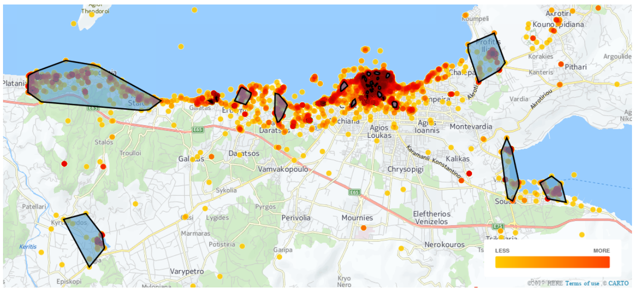

(d) Chania

Fig. 6 Extracted clusters, overlaid to the density-based visualization of photos' distributions, for the 4 selected cities.

Table 3 Performance of Xenia against the applied baseline techniques according to various metrics and for the 4 examined cities (P: precision, R: recall, best result per city is depicted in bold.)

	Athens		**Thessaloniki**		**Chania**		**Heraklion**	
	R	**P**	**R**	**P**	**R**	**P**	**R**	**P**
Xenia	**.676**	**.591**	**.717**	**.688**	**.753**	**.653**	**.779**	**.677**
GPOI	.434	.428	.551	.590	.600	.615	.623	.653
NPOI	.488	.576	.614	.647	.563	.628	.517	.568
RPOI	.462	.436	.389	.316	.388	.462	.468	.476

- *Recall* measures the fraction of retrieved documents from a collection that are relevant, in the context of a given query. Herein, we calculate Recall as the amount of visited POIs from a users travel history that also appear in the recommended route.
- *Precision* is defined as the fraction of POIs that appear in a recommended travel route and are also part of a users real-life travel substream.

In Table3 we present the results of the evaluation of Xenia in the 4 examined Greek cities. Our proposed methodology managed to outperform all the baselines that were evaluated, essentially fulfilling its goal, i.e, to maximize the overall travel experience for a tourist, given a area that is characterized by a high-density of POIs. In particular, for the Recall metric it achieved the best results in comparison with the other baselines. Similarly, for the Precision metric, Xenia was again the best. Overall, the proposed system is able to construct routes that are similar to the ones that the users followed during their vacations.

5.3.2 Qualitative Evaluation

For a qualitative evaluation of the proposed system we chose to perform a study, focusing on user satisfaction, following the paradigm of De Choudhury et al. [8] who used workers from the Amazon Mechanical Turk[17] crowdsourcing intelligence service and of [46] who also opted for an evaluation with real-life users. Since our goal is to provide travel routes that adhere to a multitude of constraints and parameters, we feel that apart from the quantitative evaluation that has been presented in subsection 5.3.1, an appropriate, complementary way of evaluation is to assess the satisfaction of real tourists.

We should emphasize that in general, the appraisal of tasks aiming at users' satisfaction is known to be a difficult and expensive procedure, which may involve empirical issues in the process [28]. However, we have conducted a user-centered evaluation by involving 15 real-life users, who although not being tourists *while* performing the evaluation, they were familiar to a great extent with the examined cities having visited them as tourists *in the past.* Users were students from two (2) academic institutions.[18]

[17] `https://www.mturk.com/mturk/welcome`

[18] More specifically we used 8 students from the Technological Educational Institute of Central Greece, Lamia, Greece and 7 students from the Ionian University, Corfu, Greece.

We did not intent to bias the opinion of the users, thus preferred not to present them with ground truth routes (i.e., we skipped the first part of the evaluation procedure proposed by [8]). Instead, we independently assessed the utility of each route, based on the users' experience on each city. The evaluation procedure was as follows: firstly we showed them a route (i.e., the set of POIs to be visited along with the recommended visit durations and finally the corresponding transition times) and then we asked them to answer the set of questions depicted in Table 4. We repeated this procedure for each route of those selected for the evaluation. We selected 6 routes from Athens, 3 routes from Thessaloniki, 2 routes from Chania and 2 from Heraklion. Since not all users had visited each city, our evaluation was based on 10 users for Athens, 6 for Thessaloniki, 5 for Chania and Heraklion.

As it can be seen in Table 4, Q1 and Q2 aim to evaluate the usefulness of a route and of the POIs that are included within it. We feel that these are somehow dependent, but a route that overall is characterized as useful, in terms of POIs visited may not satisfy the users since they may expect e.g., to visit more POIs within the given time, or more popular POIs than those contained within. Q3 aims to evaluate the satisfaction of a user, regarding the suggested time spent at the visited POIs. As inappropriate, we denote a suggested amount of time that is far more or far less than what the users feel as adequate so as to comfortably visit all the recommended POIs. Q4 evaluates the users' opinion on the walking length of a POI. Q5 aims to evaluate whether a suggested route is likely to be suggested also by a user. Q6–7 aim to assess the quality of the route in terms of view and culture.

These questions have 5 possible answers, with one being neutral, two positive and two negative, based on the good practice suggested by [22] for user evaluation, i.e., the number of positive answers should be equal to the one of negative and also a neutral answer should be provided. Finally, Q8–Q11 have been selected so as to collect valuable input from the users, concerning details of the proposed route, whereas Q12 investigates whether Xenia would be a useful mobile application.

Answer to questions 1–6 and 12 was mandatory. Thus, an appropriate metric for the answers of these questions is the mean response, which is the average from all answers. Corresponding results are depicted in Table 5. Given these results, we may argue that users that participated in the evaluation process were in general satisfied by Xenia. Best results were achieved in Athens and worst in Heraklion. It is our belief that the significantly larger number of POIs in Athens, and of course their importance played a crucial role towards these results. Moreover, the larger number of tourists in Athens provided a more consistent foundation towards the system's robustness. However, in all cities results were satisfactory, since the mean answers were in general positive while being neutral in only a few cases. Furthermore, the mean response for Q12 was 4.27, which indicates that overall the users would use a mobile application based on Xenia.

As for questions 8–11, since the answer was not mandatory, we chose to adopt the metric used by [8], namely *Mean Average Error Fraction* (MAEF), which is calculated as the average of the Mean Error Fractions, defined as

$$\mathrm{MEF}(\mathcal{R}) = \frac{1}{|U(\mathcal{R})|} \sum_{u \in U(\mathcal{I})} \frac{b(u)}{|\mathcal{R}|} \,, \tag{8}$$

where $\mathcal{R}$ is a route produced by our system and $U(\mathcal{R})$ is the set of the users that have provided a response concerning $\mathcal{R}$, while $b(u)$ and depending to the question is the number of POIs (Q8) that have been characterized as irrelevant or missing (Q9), the number of transit times the user feels are too long (Q10) or the number of visit times the user feels that are either insufficient or too long (Q11). Corresponding results are depicted in Table 6. We may observe that results are quite similar in Athens, Thessaloniki and Chania. However, in Heraklion, the users feel that quite a few points were irrelevant, many transit times were longer than what they preferred and many visit times were inappropriate. This can be attributed to the use of the whole prefecture as the bounding box of the urban area, instead of a city (i.e., as in Athens and Thessaloniki). Moreover, Chania and Heraklion have significantly less POIs that are of general interest, contrary to Athens, thus there were multiple AOIs containing a single POI, thereby leading to the detection of sparser and secluded areas, which directly affects the transit times between the recommended POIs. Also, the scores of POIs are extracted based on their popularity. However, at the evaluation phase, users are asked to rate POIs, which obviously is based upon different criteria than their popularity.

6 Conclusions and Future Work

In this paper, we presented *Xenia*, a novel system for recommending travel routes that satisfy a multitude of different constraints and parameters. We used an unsupervised approach based on HDBSCAN in order to cluster a set of geo-tagged images collected from Flickr as a means to discover high popularity areas that we defined as AOIs. As we have already justified, each AOI may contain multiple POIs that are favored by tourists during their vacations.

We then modeled the TTDP bearing in mind that a tourist typically has a limited time budget, while she/he wishes to visit as many POIs as possible. Under the assumption that a POI's popularity is proportional to the gain a tourist has upon visiting and using the Flickr users' history we then solved the TTDP problem through the integer programming formulation of the OP. We evaluated the proposed system against a set of typical baselines derived from the state-of-the-art and using 4 Greek cities, known for being popular tourist destinations. Apart from a quantitative evaluation, we also included a user evaluation aiming to measure users' satisfaction on the proposed travel routes.

We demonstrated the robustness and the efficiency of our system, considering all the necessary stages, i.e., from data clustering to travel route construction and also its feasibility, by providing all necessary implementation details.

Table 4 The questionnaire that has been completed by users for each route they were shown, during their evaluation of Xenia

Question	**Possible Answers**
Q1: Please provide your overall opinion on the usefuleness of the presented route for a tourist.	**1**: Absolutely unuseful **2**: Somehow unuseful **3**: Not useful, neither unuseful **4**: Very useful **5**: Absolutely useful
Q2: Please provide your overall opinion on the POIs contained within the presented route.	**1**: Absolutely insatisfactory **2**: Somehow insatisfactory **3**: Not satisfactory, neither insatisfactory **4**: Very satisfactory **5**: Absolutely satisfactory
Q3: Please rate the visit times suggested for each POI of the presented route.	**1**: Absolutely inappropriate **2**: Somehow inappropriate **3**: Not appropriate, neither inappropriate **4**: Very appropriate **5**: Absolutely appropriate
Q4: Please provide your opinion on the length (in km) of the presented route.	**1**: Absolutely insatisfactory **2**: Somehow insatisfactory **3**: Not satisfactory, neither insatisfactory **4**: Very satisfactory **5**: Absolutely satisfactory
Q5: Do you feel that this route successfully captures your notion of a recommended route to a tourist, given the available time?	**1**: Certainly not at all **2**: Certainly not, though it contains a few POIs I would recommend **3**: It contains many POIs I would recommend, but I would not recommend the whole route **4**: It is very close to what I would recommend **5**: I would recommend this route, as is
Q6: Do you feel that the route adequately captures the "scenic" view of the visited city?	**1**: Absolutely not **2**: Probably not **3**: It does at some point **4**: Probably yes **5**: Certainly yes
Q7: Do you feel that the route adequately captures the historical significance, traditions and culture of the visited city?	**1**: Absolutely not **2**: Probably not **3**: It does at some point **4**: Probably yes **5**: Certainly yes
Q8: Please indicate us any POIs you think that are irrelevant to the goals of a tourist	free answer
Q9: Please indicate us any POIs that you feel that are missing from the system	free answer
Q10: Please indicate us which transit times you feel are too long	free answer
Q11: Please indicate us which visit times you feel that are either insufficient or more than sufficient	free answer
Q12: Please provide whether you intend to use Xenia, provided that it evolves to a publicly available mobile app.	**1**: Absolutely not **2**: Probably not **3**: I could try it if necessary **4**: Probably yes **5**: Absolutely yes

Table 5 Mean response for Q1–Q7, for all 4 cities.

	Athens	Thessaloniki	Chania	Heraklion	Average
Q1	4.4	3.8	4.1	3.4	3.9
Q2	4.6	3.6	3.6	3.9	3.9
Q3	4.3	3.8	4.4	3.0	3.9
Q4	4.2	4.1	4.2	3.6	4.0
Q5	4.6	4.0	3.8	3.2	3.9
Q6	4.3	4.3	4.5	4.4	4.4
Q7	3.9	4.3	4.3	4.1	4.2

Table 6 Fraction of irrelevant POIs (Q8), missing POIs (Q9), too long transit times between POIs (Q10) and inappropriate visit times for POIs (Q11), for all 4 cities.

	Athens	Thessaloniki	Chania	Heraklion	Average
Q8	0.11	0.17	0.19	0.73	0.30
Q9	0.43	0.42	0.40	0.30	0.39
Q10	0.08	0.19	0.11	0.26	0.16
Q11	0.13	0.26	0.15	0.37	0.23

Since no prior knowledge was necessary at any stage, we also proved that our approach is able to successfully work "in the wild." Our quantitative results indicate that Xenia is able to provide routes that are comparable to real-life ones. Furthermore, the qualitative results showcase the participants' positive response regarding both the recommended POIs and the suggested visit times, respectively, an observation which was in general consistent among all 4 examined cities. Moreover, the user evaluation process indicated that users were satisfied with the proposed routes in the majority of cases.

Of course, the evaluation results could be further improved. Among our immediate plans for future work, we intend to assess our proposed system in additional cities and integrate temporal and/or weather data as a means to deduce the number of people that would probably visit a particular POI at a given time. Focus will be given on personalization by providing personalized travel routes upon detection of user preferences and the construction of themed routes e.g., recommendation of POIs deriving solely from a particular archaeological period, such as classical Athens, byzantine period etc. We also plan to make suggestions for groups of people, instead of individuals. Moreover, we would like to expand the domain of recommendations by adding other types of POIs, such as restaurants and bars, by integrating the ability to schedule breaks during the trip for food and/or coffee etc. Also, we would like to provide alternative means of transportation such as bikes, buses, trams and the metro.

Further research directions that could be followed include the introduction of semantics in the geo-clustering process, the automatic AOI characterization (e.g., the identification of archaeological, commercial, nightlife districts etc.) and consequently the detection of the most prevalent POI category within an AOI. Within the process, information from other social networks, e.g., Twitter,[19] blogs etc. and other modalities, e.g., sentiment analysis using Natural Language Processing may be incorporated. Overall, we feel that within the next few years, the research area of personalized trip suggestions will continue its growth and further research disciplines shall contribute towards more efficient and fully automatic personalized solutions. We also feel that crucial shall be the role of a huge dataset made publicly available from Yahoo! and consisting of 100M Flickr images, accompanied by their metadata [43].

References

1. Y. Arase, X. Xie, T. Hara and S. Nishio, *Mining people's trips from large scale geo-tagged photos*, In Proc. of ACM Int'l Conf. on Multimedia (MM), 2010.
2. I. Brilhante, J.A. Macedo, F.M. Nardini, R. Perego and C. Renso, *Where shall we go today?: planning touristic tours with Tripbuilder.* In Proc. of ACM Int'l Conf. on Information & Knowledge Management, 2013.
3. R.J. Campello, D. Moulavi and J. Sander, *Density-based clustering based on hierarchical density estimates.* In Proc. of Pacific-Asia Conf. on Knowledge Discovery and Data Mining, 2013
4. L. Cao, J. Luo, A. Gallagher, X. Jin, J. Han, and T.S. Huang, *A worldwide tourism recommendation system based on geotagged web photos.* In Proc. of IEEE ICASSP, 2010.
5. R. Chalfen, *Snapshot versions of life.* University of Wisconsin Press, 1987.
6. I. Chao, B. Golden, E. Wasil, *Theory and methodology a fast and effective heuristic for the orienteering problem.* European Journal of Operational Research 88, 475489, 1996
7. Y.Y. Chen, A.J. Cheng and W.H. Hsu, *Travel recommendation by mining people attributes and travel group types from community-contributed photos.* IEEE Trans. on Multimedia, 15(6), pp.1283–1295, 2013.
8. M. De Choudhury, M. Feldman, S. Amer-Yahia, N. Golbandi, R. Lempel and C. Yu. *Automatic construction of travel itineraries using social breadcrumbs.* In Proc. of ACM Conf. on Hypertext and hypermedia, 2010.
9. Deloitte. *Photo sharing: trillions and rising*, `https://www2.deloitte.com/content/dam/Deloitte/global/Documents/Technology-Media-Telecommunications/gx-tmt-prediction-online-photo-sharing.pdf` (retrieved 2017-01-17)
10. M. Duggan, *Photo and Video Sharing Grow Online*, `http://www.pewinternet.org/2013/10/28/photo-and-video-sharing-grow-online/` (retrieved 2016-11-16).
11. M. Ester, H.P. Kriegel, J. Sander and X. Xu, *A density-based algorithm for discovering clusters in large spatial databases with noise.* In KDD, vol.96, no.34, pp. 226–231, 1996.
12. D. Gavalas and M. Kenteris, *A web-based pervasive recommendation system for mobile tourist guides.* Personal and Ubiquitous Computing, 15(7), pp.759–770, 2011.
13. A. Gionis, T. Lappas, K. Pelechrinis and E. Terzi, *Customized tour recommendations in urban areas.* In Proc. of ACM Int'l Conf. on Web Search and Data Mining, ACM, 2014.
14. F. Girardin, F. Calabrese, F. Dal Fiore, C. Ratti and J. Blat, *Digital footprinting: Uncovering tourists with user-generated content*, IEEE Pervasive Computing 7(4), pp. 36-43, 2008.
15. S. Greenwood, A. Perrin and M. Duggan, `http://www.pewinternet.org/2016/11/11/social-media-update-2016/` (retrieved 2017-01-17)
16. Q. Hao, R. Cai, J.-M. Yang, R. Xiao, L. Liu, S. Wang and L. Zhang, *Travelscope: standing on the shoulders of dedicated travelers.* In Proc. of ACM MM, 2009.
17. L. Hollenstein and R. Purves, *Exploring place through user-generated content: Using Flickr tags to describe city cores*, Journal of Spatial Information Science, no. 1 pp.21–48, 2010.

[19] `http://www.twitter.com`

18. H.P. Hsieh, C.T. Li and S.D. Lin, *Measuring and recommending time-sensitive routes from location-based data*. ACM Trans. on Intelligent Systems and Technology (TIST), 5(3), p.45, 2014.
19. Y. Hu, S. Gao, K. Janowicz, B. Yu, W. Li and S. Prasad, *Extracting and understanding urban areas of interest using geotagged photos*, Computers, Environment and Urban Systems, 54, pp.240–254, Elsevier, 2015.
20. S. Jain, S. Seufert and S. Bedathur, *Antourage: Mining Distance-Constrained Trips from Flickr*. In Proc. of ACM WWW, 2010.
21. S. Jiang, X. Qian, J. Shen and T. Mei *Travel Recommendation via Author Topic Model Based Collaborative Filtering*. Lecture Notes in Computer Science, Vol. 8936, pp.392–402, Springer, 2015.
22. R. Johns, *Likert items and scales*, 2010, `https://www.ukdataservice.ac.uk/media/262829/discover_likertfactsheet.pdf` (retrieved 2017-01-31)
23. S. Kisilevich, D. Keim, N. Andrienko and G. Andrienko, *Towards Acquisition of Semantics of Places and Events by Multi-perspective Analysis of Geotagged Photo Collections*. Geospatial Visualisation, pp. 211–233, Springer, 2013.
24. T. Kurashima, T. Iwata, G. Irie and K. Fujimura, *Travel route recommendation using geotags in photo sharing sites*. In Proc. of ACM Int'l Conf. on Information and Knowledge Management, 2010.
25. K.H. Lim, J. Chan, C. Leckie and S. Karunasekera. *Personalized Tour Recommendation Based on User Interests and Points of Interest Visit Durations*. In Proc. of Int'l Joint Conference on Artificial Intelligence (IJCAI), 2015.
26. K.H. Lim, J. Chan, C. Leckie and S. Karunasekera, *Towards Next Generation Touring: Personalized Group Tours*. In Proc. of ICAPS, 2016.
27. V.I. Levenshtein, *Binary codes capable of correcting deletions, insertions, and reversals*, Soviet Physics Doklady 10, pp. 70710, 1966.
28. Y. Liu, J. Bian and E. Agichtein, *Predicting information seeker satisfaction in community question answering*, in Proc of int'l ACM SIGIR conf. on Research and Development in Information Retrieval, 2008.
29. J. Liu, Z. Huang, L. Chen, H.T. Shen and Z. Yan, *Discovering areas of interest with geo-tagged images and check-ins*, Proc. of ACM Int'l Conf. on Multimedia (MM), 2012.
30. X. Lu, C. Wang, J.M. Yang, Y. Pang and L. Zhang, L., *Photo2trip: generating travel routes from geo-tagged photos for trip planning*. In Proc. of ACM Int'l Conf. on Multimedia (MM), 2010.
31. E.H.C. Lu, C.Y. Chen and V.S. Tseng, *Personalized trip recommendation with multiple constraints by mining user check-in behaviors*. In Proc of Int'l Conf. on Advances in GIS, ACM, 2012.
32. E.H.C Lu, S.H. Fang and V.S. Tseng, *Integrating tourist packages and tourist attractions for personalized trip planning based on travel constraints*. GeoInformatica, 20(4), pp.741–763, 2016.
33. A. Majid, L. Chen, Ling, G. Chen, H.T. Mirza, I. Hussain and J. Woodward, *A context-aware personalized travel recommendation system based on geotagged social media data mining*. Int. Journal of Geographical Information Science, vol.27, no.4, pp. 662–684, Taylor & Francis, 2013.
34. C.E. Miller, A.W. Tucker and R.A. Zemlin, *Integer programming formulation of traveling salesman problems*. Journal of the ACM (JACM), 7(4), pp.326–329, 1960.
35. D. Quercia, R. Schifanella and L.M. Aiello, *The shortest path to happiness: Recommending beautiful, quiet, and happy routes in the city*. In Proc. of ACM Conf. on Hypertext and Social Media, 2014.
36. E. Schubert, A. Koos, T. Emrich, A. Züfle, K.A. Schmid, A. Zimek, A Framework for Clustering Uncertain Data. Proc. of the VLDB Endowment, 8(12): pp.1976–1979, 2015.
37. S. Martello, and P. Toth. *Knapsack problems: algorithms and computer implementations*. John Wiley & Sons, Inc., 1990.
38. A. Popescu and G. Grefenstette, *Deducing trip related information from flickr*. In Proc. of ACM WWW, 2009.
39. A. Popescu, G. Grefenstette and P.-A. Moëllic, *Mining tourist information from user-supplied collections*. In Proc. of ACM CIKM, 2009.

40. W. Souffriau, P. Vansteenwegen, J. Vertommen, G. Vanden Berghe and D. Van Oudheusden, *A personalized tourist trip design algorithm for mobile tourist guides*. Applied Artificial Intelligence, vol.22, no.10, pp.964–985, 2008.
41. E. Spyrou and Ph. Mylonas, *A survey on Flickr multimedia research challenges*. Engineering Applications of Artificial Intelligence, vol. 51, pp. 71-91, Elsevier, May 2016.
42. Y. Sun, H. Fan, M. Bakillah and A. Zipf, *Road-based travel recommendation using geo-tagged images*. Computers, Environment and Urban Systems, Elsevier, 2013.
43. B. Thomee, D.A. Shamma, G. Friedland, B. Elizalde, K. Ni, D. Poland, D. Borth and L. Li, *YFCC100M: The New Data in Multimedia Research*, Communications of the ACM, 59(2), pp. 64-73, 2016.
44. S. Van Canneyt, S. Schockaert, O. Van Laere and B. Dhoedt, Bart, *Time-dependent recommendation of tourist attractions using Flickr*. In Proc. of BNAIC, 2011.
45. N.A. Van House, *Personal photography, digital technologies and the uses of the visual*. Visual Studies, 26(2), pp.125–134, 2011.
46. P. Vansteenwegen, W. Souffriau, G.V. Berghe and D. Van Oudheusden, *The city trip planner: an expert system for tourists*. Expert Systems with Applications, 38(6), pp.6540–6546, 2011.
47. P. Vansteenwegen, W. Souffriau, D. Van Oudheusden, *The orienteering problem: A survey*. European J. of Oper. Research, vol.209, no.1, 2011.
48. B. Wu, Y. Murata, N. Shibata, K. Yasumoto and M. Ito, *A method for composing tour schedules adaptive to weather change*. In Intelligent Vehicles Symposium, IEEE, 2009.
49. A. Yahi, A. Chassang, L. Raynaud, H. Duthil and D.H.P. Chau, *Aurigo: an interactive tour planner for personalized itineraries*. In Proc. of Int'l Conf. on Intelligent User Interfaces, ACM, 2015.
50. H. Yoon, Y. Zheng, X. Xie and W. Woo, *Social itinerary recommendation from user-generated digital trails*. Personal and Ubiquitous Computing, 16(5), pp.469-484, 2012.