Pers Ubiquit Comput manuscript No.
(will be inserted by the editor)

Using Gamification to Discover Cultural Heritage Locations from Geo-tagged Photos

Armir Bujari · Matteo Ciman · Ombretta Gaggi · Claudio E. Palazzi

Received: date / Accepted: date

Abstract Many enchanting cultural heritage locations are hidden from tourists, especially when considering countries full of historic attractions. Tourists tend to consider only mainstream monuments and towns, neglecting wonderful little jewels along their travel itinerary. However, this is generally not their fault, as travelers cannot be aware of all the surrounding beauties when visiting a new region. To this aim, we discuss and analyze here *PhotoTrip*, an interactive tool able to autonomously recommend charming, even if not mainstream, cultural heritage locations along travel itineraries. *PhotoTrip* is able to identify these points of interest by gathering pictures and related information from Flickr and Wikipedia and then provide the user with suggestions and recommendations. An important technical challenge for this kind of services is the ability to provide only the most relevant pictures among the many available for any considered itinerary. To this aim, we have exploited social networks, crowdsourcing and gamification to involve users in the process of improving the response quality of our system.

Keywords Cultural Heritage · Gamification · Geolocalization · Knowledge Discovery

A. Bujari, O. Gaggi and C. E. Palazzi
Department of Mathematics
University of Padua, Italy
E-mail: {abujari, gaggi, cpalazzi}@math.unipd.it

M. Ciman
Quality of Life Group
University of Geneva, Switzerland
E-mail: Matteo.Ciman@unige.ch

1 Introduction

Over the past decades, digital technologies have been put to good use in many ways to improve how people could enjoy cultural heritage. Technology can digitally reconstruct historic artifacts, landscapes or events, broaden the accessibility of culture through remote distribution platform, let people learn interactively about culture, augment the visit to a museum, efficiently organize visitors' flows, create an Internet of Cultural Things, and much more [4] [8] [9] [10] [11] [43]. Digital practices can hence effectively enhance how visitors enjoy culture and learn from it [35].

One key challenge in this context is how to help tourists organize their visit to a new region without experiencing only mainstream attractions. Indeed, to plan a visit to a new city or region, travelers generally search for information in guidebooks, popular web sites or blogs. Unfortunately, these sources of information cannot cover all possible locations and tend to focus on well known (and crowded) ones. Personal travel blogs may be helpful in having access to a particular traveler's view, but this perspective could not match the experience of others and its originality strongly depends on the time invested in planning the trip; furthermore, searching blogs for this kind of information may result time consuming and not so easy for unexperienced users.

Particularly relevant is the case of a traveler willing to know whether other interesting, even if not widely known, attractions are present along his/her trip. Indeed, many enchanting cultural heritage locations are left out of the mainstream touristic circuit and this is especially true when considering countries that are full of history and culture. Travelers cannot be aware of all the surrounding possibilities when visiting a new

Fig. 1: Villa Pisani in Stra, Venice, Italy.

region; it is hence not their fault if they end up visiting only mainstream locations, loosing the possibility to enjoy enchanting (and not crowded) beauties along their itinerary. We call these places LPOIs, *Local Points Of Interest*. For instance, if you travel from Padua to Venice in Italy, an example of *LPOI* would be Stra, a little city shadowed by the close Venice, but with a fine villa of the venetian noble Pisani family (see *Villa Pisani* in Figure 1), which requires a deviation of less than 5 km from the original route.

To this aim, we here discuss and analyze *PhotoTrip*, an interactive tool we developed, able to autonomously recommend LPOIs along a tourist's route. These LPOIs are identified by gathering, analyzing and clustering the geo-tagged photos as well as related information and metadata available on the Web. In particular, we have used Flickr as the source of these photos since it represents the leading online service for managing and sharing pictures. Our system also allows to filter images with respect to a particular category requested by the user (e.g., nature, architecture, etc.) or the maximum deviation allowed. Furthermore, it also integrates other online services, such as Wikipedia, to enhance the information provided for the recommended local attractions. Finally, once the user has chosen the LPOIs, it calculates the route showing the directions to follow.

PhotoTrip embodies a new and original way to exploit the many geo-tagged and time-stamped photos available on the Web so that travelers' experience can be used to unveil charming local cultural attractions. It represents a service which is currently not offered by other Web alternatives. For instance, Google Maps provides pictures related to a required route, but it does not allow to calculate route deviations and does not filter images with respect to the user's preference in terms of category of interest or maximum route deviation. Furthermore, Google Maps does not recognize points of interest, it simply displays all the related pictures geo-tagged in the chosen rectangular area.

A crucial component of *PhotoTrip* is the ability to provide only the most relevant pictures among the many available on Flickr for any considered itinerary. Unfortunately, tags associated to pictures on Flickr are assigned by users and are hence highly susceptible to the user perception of the context. To tackle this limitation, our system allows end-users to refine the displayed content through feedbacks. However, this is just a partial solution as users are known to be selfish and not always trustworthy [34] [38]; for these reasons, the reliability of *PhotoTrip* cannot depend solely on their willingness to help.

As a further contribution of this work, we have hence used gamification to involve users in the process of improving the response quality of our system. Indeed, gamification has been proven to be effective in many scenarios where an incentive is required to convince users to effectively contribute to the feedback loop [19] [32] [36] [42]. In this paper, we show how this technique could be adopted to discern between relevant and non-relevant photos and analyze its efficacy in this context.

The rest of the paper is organized as follows. Section 2 discusses similar approaches in literature. Section 3 describes the system and its implementation. Our gamification approach is presented in Section 4 and its evaluation is reported in Section 5. We conclude in Section 6.

2 Related Work

Through this section we review literature material related to our problem. We organize our review by discussing each body of work separately and contrasting them, whenever possible, with our approach.

2.1 Route planning

The ever increasing amount of data produced by digital devices related to places and locations on the earth has raised the problem on how it is possible to analyze them in order to extract useful information. These data can be related to photos taken by individuals, GPS traces, or even check-ins made using platforms such as Facebook or Foursquare.

Different approaches analyze how to reconstruct and mine travel itineraries considering different online databases. For example, Crandall et al. [13] proposed a proof-of-concept system based on Flickr photos, inferring a relational structure between them. They define Points Of Interest (POIs) as places which have the higher number of photos posted, and propose the implementation of an online travel guidebook able to suggest attraction sites deserving a visit during a vacation period. De Choudhury et al. [14] reconstructed the travel paths a user has taken based on his/her posted online photos. These inferred itineraries from different users are then merged into a POI graph which is used to construct suggested travel paths. Similarly, Yamasaki et al. [49] proposed a system for inter- and intra-city travel recommendation based on geo-tagged pictured downloaded from Flickr. The system is based on users' preferences, a similarity distance measure between different users, and on seasonal/temporal information attached to the photos.

Harota et al. in [23] proposed a method to detect and visualize relations among places where many photos have been taken, called hotspots. Each hotspot reveals users' interests. The system is based on Flickr photos' locations and is able to organize hotspots detecting and assessing relations between shooting spots and photo subjects.

Arase et al. in [5] defined a set of trip themes, i.e., landscapes or cities, historical buildings or modern art, and mined frequent trip patterns for each theme category. Kurashima et al. instead considered the spare time of the user and in [27] presented a framework for travel route recommendation which is able to suggest different travel paths according to it. The system is based on the analysis of both geo-tagged and time-stamped photos of photo sharing services. Even in this case, the authors mined the users' photos to learn personal travel history using a probabilistic approach.

Kofler et al. addressed a similar problem in [26]. *Near2me* is a prototype system that recommends places which are not necessary touristic but are genuinely representative of the considered location, representing also a users' personal interests. The system allows users to explore, evaluate and understand recommendations, and discover informative support material, based on user photos from the image sharing community Flickr.

Another example is *Photo2Trip* [29], an automatic travel route planner which collects photos from Panoramio in order to suggest customized trip plans according to users' preferences. It is able to suggest the most popular destinations to visit. Differently from the other approaches, the system allows the users to interactively specify their preferences.

More complex is the approach proposed by Popescu and Grefenstette in [37]. In this case, the system suggests touristic visits based on the aggregation of photo's annotation, building a similarity matrix between users and photos in order to provide the best suggestion. GPSView [50] is a scenic driving route planner, an augmented GPS navigation system. The goal of the system is to plan a driving route with scenery and sightseeing locations, allowing travelers to enjoy sightseeing while driving. This is achieved by exploiting a database of scenic roadways under the implicit assumption that a multitude of photos taken along a roadway implies that this roadway is probably appealing and catches the public's attention. The system creates a set of POIs which have good scenic qualities and visibility, formulating scenic route planning as an optimization task toward the best tradeoff between sightseeing experience and traveling distance.

Finally, other approaches do not rely only on images, but even on more complex data sources. For example, Shen et al. [44] proposed a landmark search system based on the aggregation of heterogeneous tourist information, to provide smart travel guide, such as landmark ranking system. In this case, even user ratings are used and analyzed, to measure social popularity of each landmark, as of pictures and images collected from Flickr. *TripRec* [24] uses Foursquare and Gowalla to recommend trip routes by leveraging large-scaled check-in data through mining the moving behaviors of users. It allows to specify start/end and must-go locations, plus some flexibility parameters like the expected duration of the trip. Moreover, it uses time of check-in points to understand which are the people's frequent sequences of locations.

All the cited systems address the problem of identifying (and defining) the most popular POIs in order to create a travel path towards them. Our approach is different since the user already knows his/her destination, but he/she also wants to investigate if there exist some minor, but noteworthy, local attractions within a defined deviation limit. Moreover, our system retrieves additional information from Wikipedia to help the user's choice and uses gamification techniques to filter source images in order to provide representative results.

2.2 Gamification and serious games

Gamification and *serious games* have gathered a lot of interest in the recent years. Gamification is defined as the usage of game elements, e.g., badges, rewards, competition, leaderboards, etc., in non-gaming systems, making individuals more engaged in completing a specific task or reaching a goal [1] [17] [16].

As an example, in [47] the authors proposed a gamification solution providing incentives to people in participating in heavy physical and mental sensing tasks, i.e., they asked users to check-in at specific places where the number of places was much higher than the number of users. The solution adopts a status level, earned points, ranking scheme and badges to motivate people. The use of gamification resulted in a 20% increase of user participation.

An example related to our work is presented in [45], where the authors describe and implement a gamification approach for image sensing and tagging. In particular, they adopt gamification techniques to make image annotation more engaging, since this task is still very cumbersome for most of the users. Differently from our approach they ask users to annotate images rather than ask people about the correctness of the annotation; they aim at finding a balanced agreement rate between pre-established annotations and those defined by the users. Mekler et al. in [31] discussed the effects of points and meaningful framing on intrinsic motivation and performance in an image annotation task. They asked users to provide tags that describe the mood of 15 abstract paintings. They observed how points elicited more tags from participants, since they establish a connection between user effort and performance. Moreover, meaningful frames seemed to inspire participants to work harder and take their time when labeling images, increasing the overall quality of the provided feedback. Even in this approach, the authors ask users to annotate an image and not to vote on the correctness of an annotation.

Differently from gamification, the idea at the basis of *serious games* is to teach something serious, useful or healthier, in a funny way, introducing not only game elements like gamification, but even game play and a real game in this teaching effort [32] [39]. In this way, users, while playing the game, can learn a specific task [46], interiorize a healthier behavior [3], or follow a medical rehabilitation program [40] [21], [15], etc.

Bellotti et al. [6] presented the SandBox Serious Game, a framework for *serious games* development in cultural heritage which relies on a generalization of task-based learning theory. The model invites players to perform cognitive tasks while exploring information-rich virtual environments. Results showed how games appeared particularly suited for supporting the study of images, especially iconography. Compared to reading text, a game forces the player to focus more intensively on problems, which favors knowledge acquisition and retention.

MuseUs [12] is a pervasive *serious game* for museums; players are invited to create their own exposition with the aim of providing a learning effect during the visit. Moreover, the application stimulates the visitor to look at cultural heritage elements in a different way, permitting the construction of personal narratives while creating a personal exposition.

Similar, Marfia et al. [41] [30] have proposed a serious game able to preserve and share handcraftship cultural heritage. The system is based on capturing and recognizing the gestures performed by the user's hands through a camera and follows a serious game approach in verifying the ability of the user in performing the correct set of movements an artisan would do. This kind of systems has demonstrated to be effective in preserving and passing along handcraftship skills and cultural heritage in general [33].

O'Munaciedd [28] is a *serious game* for children for improving concentration, memorization and problem solving abilities, providing information about the artistic and cultural heritage of the city of Matera. The application is based on a series of 11 mini-games that appear whenever a child approaches a historical site. Data analysis has shown that the game was enjoyable for children, demonstrating the potential of the approach and social networks in the construction of knowledge.

H-Treasure Hunt [25] integrates a location-based service with object based sensors for artifacts at a historic place. In this game, players wear a head mounted display and explore historic sites interacting with artifacts to complete missions. The overall objective of the game is to support cultural heritage teaching and learning, enhancing visits. Users agreed that the play element proved to be an appealing factor and that the gesture and voice interaction made the exploring experience more stimulating.

Froschaugen et al. [18] proposed *ThIATRO*, a multiplayer online *serious game* that helps students in learning art history. The system combines aspects of learning and fun, creating a game-like environment, raising interest in art history, contemporary culture and cultural heritage. The evaluation of the system showed its ability to change the players' aesthetic responses, allowing them to perceive art on a deeper level than prior to game-play. Moreover, it has the potential to continuously motivate a younger generation to deal with culture on a user-friendly plane.

As we can see, gamification and *serious games* techniques are very effective in engaging users and have the potential to involve them in tasks that otherwise would be considered boring, e.g., image ranking/labeling.

3 System Description

Through this section we provide a functional description of the system components, outlining their salient features and *modus operandi*. The gamification aspect, the approach taken in order to acquire user feedback and improve the system efficacy, is presented in Section 4.

To start a search, the user needs to fill in the data form shown in Figure 2. The mandatory input includes a start point and a destination[1] for the travel, the maximum deviation allowed from the main course, the travel means (by car or by feet) and the user's interests[2].

Once these search criteria are specified and a search is issued through the web interface, the following processing steps follow.

1. The web browser checks that the input data provided by the user are consistent (e.g., a radial search does not contain a destination point). If the data are correct, they are sent to the server, which calculates the path and the search area.
2. To minimize the waiting time, the server-side logic implements a *divide et impera* approach: the search area for LPOIs is divided into a grid of boxes, where each box is identified by two points, a and b, representing, respectively, the upper-left and the right-bottom corners of the box. Each box is associated to a sub-quest of the original problem, i.e., a search into a smaller area. Then, for each sub-quest, the communication between client and server proceeds as follows:
   (a) the client sends a Flickr request to the server for the proper box;
   (b) the server retrieves the results on behalf of the client and sends back a response;
   (c) the client fills a photo gallery with the received results. Each sub-quest contributes to the final result, depicted in Figure 3, if a LPOI is found. The user can play a presentation of the pictures regarding LPOIs found so far, or can choose a single photo;
   (d) the route is complete when all the sub-quests have been fulfilled.
3. At the end of the computation, the system notifies the user (see Figure 3).

All the sub-quests are independent from each other; they can hence be executed in parallel.

The approach of dividing the search into smaller sub problems was intentional. The rational is to provide the user with an incremental outcome without him/her waiting for the complete response. Step 2c exemplifies our choice, allowing the user to interact with a set of pictures and LPOIs as soon as the computation of the first search box has ended or the first LPOI is found. Then, incrementally, the server introduces other LPOI photos until all the search boxes are filled in.

[1] The radial search requires only the start point and searches within a maximum radius.

[2] The system is equipped with some predefined categories, e.g., architecture (containing the keywords palace, building, villa, castle, etc), nature (containing nature, tree, mountain, etc), recent photos (simply selecting the newest photos), etc.

3.1 Problem definition

Before delving into the technicalities of *PhotoTrip*, we first need to discuss the data source our system relies upon. *PhotoTrip* mines collections of pictures from Flickr, filtering them with respect to a set of criteria. From the system's perspective, a collection of photos is a set $C = \{p_i\}$ where each p_i is a list

$$\langle id_{p_i}, geo_{p_i}, T_{p_i}, fav_{p_i}, v_{p_i}, c_{p_i} \rangle$$

where:

- id_{p_i} denotes a unique identifier;
- $geo_{p_i} = \langle lt_{p_i}, lg_{p_i} \rangle$ denotes the location where the picture was taken and represents the latitude (lt_{p_i}) and the longitude (lg_{p_i}) of the photo's location;
- $T_{p_i} = \{tag_j\}$ denotes a set of tags associated with the photo p_i;
- $fav_{p_i} \in \mathbb{N}$ denotes the number of users who have added p_i to the set of their favorite photos;
- $v_{p_i} \in \mathbb{N}$ denotes the number of visualizations for picture p_i
- $c_{p_i} \in \mathbb{N}$ denotes the number of comments posted for p_i.

Flickr is a large database of photos and not all of them may be relevant to the quest at hand. Therefore, the system is required to filter the fetched content according to the users' search criteria. Indeed, in order to retrieve relevant photos from Flickr, the system compares each T_{p_i} with a set of tags T_u defined by the user through the web interface. *PhotoTrip* allows the user to choose from a set of predefined categories (e.g., architecture, nature, etc.); in alternative, the user can provide a custom set of tags. For each category, the system has a built-in set of tags, computed *a priori* or based on past analysis of the Flickr data set.

In specifics, the photos are retrieved on the basis of location data (geo_{p_i}) and the set T_{p_i} provided by the user that has taken the photo. The obtained collection of photos C, representing photos within the bounding box accounting for the maximum deviation defined by the user, is analyzed and filtered to retrieve a list of LPOIs for each box in the route. To this end, the first step is to filter the collection C to select only relevant pictures, then the remaining pictures are clustered into LPOIs.

Once the photos have been associated to their respective boxes, depending on the set cardinality, an additional filtering step might be necessary. An approach to this filtering step would be to consider a combined metric considering the number of times a user has marked a particular photo as his/her favorite (fav_{p_i}), how many times a photo was visualized (v_{p_i}) and the

Fig. 2: Screenshot of *PhotoTrip*: the initial form.

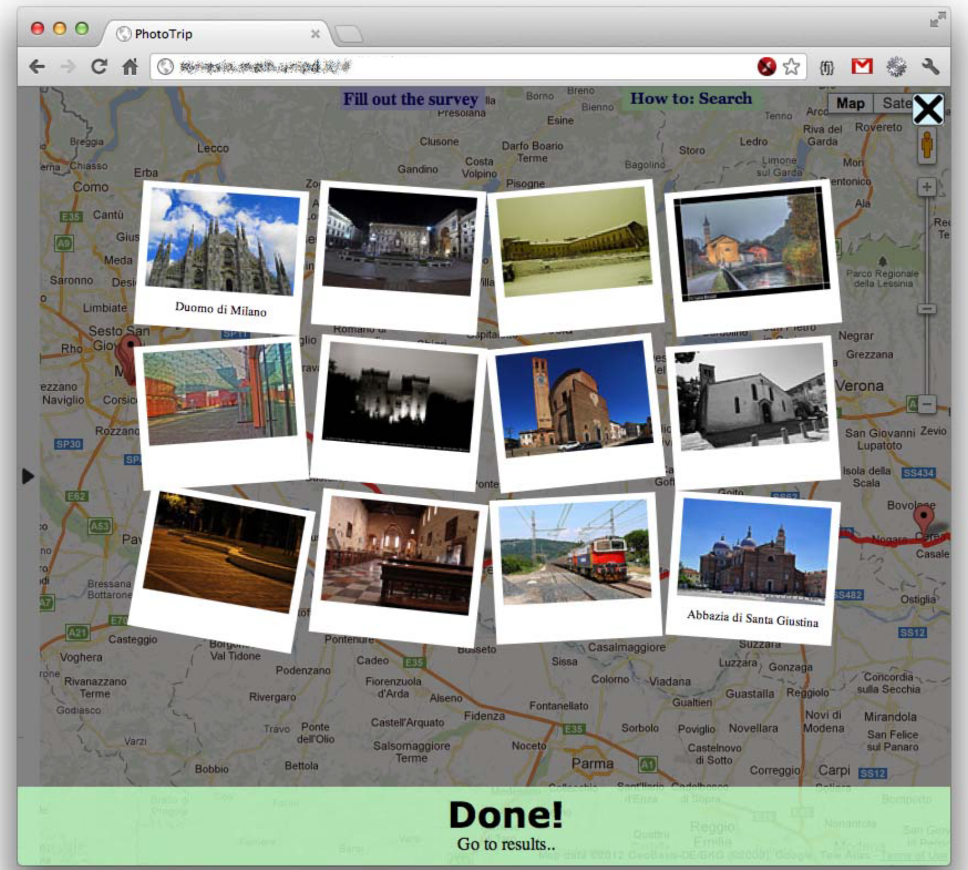

Fig. 3: Screenshot of *PhotoTrip*: LPOI gallery.

number of comments the photo has received (c_{p_i}). This would produce a value that needs to be greater than a threshold, which is empirically calculated analyzing the photo data set from the Flickr community.

However, this approach gives rise to the following limits which are inherent to the process of acquiring this additional information. Indeed, in order to obtain the information regarding the number of users who have added a particular photo to his/her favorites, an additional request to the Flickr API is required. Therefore, additional resources are needed to handle this *modus operandi*. Similarly, in order to obtain the number of comments of a photograph, additional extra requests and extra time are needed. We must recall that, unlike other recommender systems for touristic travel, our goal was to develop an online system relying on Flickr as a data source. Unfortunately, the Flickr API has constraints, granting users a maximum of one request per second. Hence, a filtering scheme based on fav_{p_i} and c_{p_i} requires, at least, two additional seconds for each analyzed picture, i.e., an excessive increase in processing time.

More importantly, the number of favorites and comments of a picture might not be a good indicator of the relevance of a picture to that LPOI. It might be the case of a picture related to a particular user or community, having no relevance to the quest at hand. This information might be counterproductive, if not harmful, for our system.

Therefore, we have decided to omit the analysis of the frequency of comments and the number of times that a picture has been marked as favorite; instead, we have developed a filter based on the number of visits received by each photo. In this scheme, a picture is considered as relevant when the metric is above a certain threshold and the threshold is computed dynamically for each set of pictures so as to always guarantee a list of LPOIs being returned. In this approach, the time needed for each search is lowered while maintaining a good quality and quantity of the photos displayed by *PhotoTrip*.

At this point, the retrieval and selection of the pictures is complete. As already discussed through this section, the photo selection process is performed for each box within the original bounding area. Each of these boxes is identified by two points, a and b, representing, the upper-left and the right-bottom corners of the box, respectively. Given the set $box_{a,b}$, which represents the infinite set of points contained in the box identified by a and b and a set of tags T_u given as input by the user, we can define the obtained collection of pictures as:

$$\begin{aligned} C^f_{a,b} \;=\; \{p_i|\; & geo_{p_i} \in box_{a,b} \quad \wedge \\ & \exists j | tag_j \in T_{p_i} \cap T_u \wedge \\ & v_{p_i} > threshold_{a,b} \;\} \end{aligned}$$

where the $threshold_{a,b}$ is computed as the arithmetic mean of the visits received by each photos in $box_{a,b}$.

For each sub-quest, if the filtered collection $C^f_{a,b}$ is not empty, the next step is to aggregate this pictures into LPOIs. To this end, the system aggregates pictures according to their position on the map. All the pictures with a distance less than a threshold are considered as belonging to the same LPOI. A LPOI contained in $box_{a,b}$ can be defined as a tuple:

$$LPOI = (lt_{p_c}, lg_{p_c})$$

where lt_{p_c} and lg_{p_c} are, respectively, the latitude and the longitude of a picture chosen as geo-tag for that particular LPOI. Given a candidate photo p_c for a LPOI and a threshold distance δ, the set of pictures contained is defined as:

$$\{p_i | p_i \in C^f_{a,b} \wedge |p_c - p_i| \leqslant \delta \;\}$$

where δ is calculated as the minimum value between the maximum deviation defined by the user and 400 m for a travel by car and 100 m for a travel by foot[3]. The candidate photo p_c is the photo depicting the LPOI with the lowest value of longitude.

In specifics, the set of images is aggregated into LPOIs as follows:

1. the set of photos is sorted according to their longitude;
2. the first picture is removed from the set and is considered as a candidate picture p_c for a new LPOI;
3. all the photos with distance less then δ from p_c are removed from the set and added to the new LPOI;
4. step 2 and 3 are repeated until the initial set of photos is empty.

To avoid ambiguities, we need to point out that the candidate photo for an LPOI p_c is not necessarily a representative photo for that particular LPOI from a semantic point of view and it has no particular relevance in the visualization of the LPOIs, where all the pictures are displayed in a balloon (see Figure 4). It is simply a starting point for the computation.

Concluding, the set of LPOIs returned to the user is the union of all LPOIs found so far.

[3] These values for δ have been stated after an analysis of images of the Flickr community.

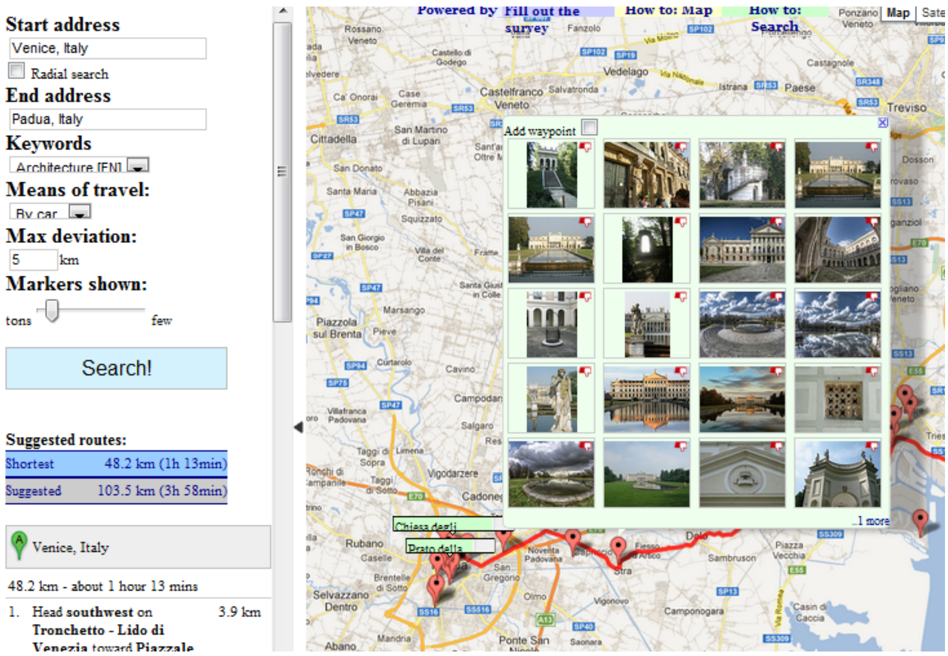

Fig. 4: Screenshot of *PhotoTrip* regarding Villa Pisani.

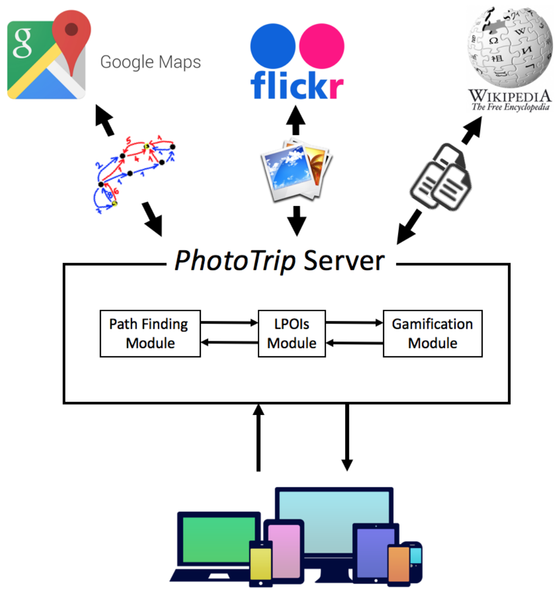

Fig. 5: The architecture of the *PhotoTrip* system.

3.2 System implementation

The architecture of the system is depicted in Figure 5. Addressing system efficiency and scalability, whenever possible, *PhotoTrip* shifts some data processing to the client side while the server-side is exploited to retrieve photos from Flickr, manage cached routes and additional information (i.e., link to Wikipedia). Addressing the resource constrained nature of mobile devices, we are currently involved in a mobile version of the client-side software where the computational burden is entirely shifted to the cloud.

Our system exploits different multimedia services available on the web, i.e., Google Maps, Flickr and Wikipedia API [48]. Each time a travel route search is performed, the system goes through the following steps:

1. it exploits Google Maps to search for the shortest path from the input starting point to the input destination point;

2. it computes the partitions of the overall bounding area into smaller boxes;
3. for each of the computed box, it retrieves and filters photos from the Flickr community and fills in the photo gallery;
4. it automatically retrieves necessary information related to what is displayed in the photos from Wikipedia;
5. it computes and displays the suggested path.

The partitioning of the bounding box denoted by the travel route is computed using a customized version of RouteBoxer [22], a JavaScript library developed by Google with the aim of managing travel paths. This box is divided into a grid where each squared cell has the side equal to the maximum deviation accepted by the user. A second step marks the cells which contain a route segment and the adjacent ones as the area to search for LPOIs.

The system thereafter combines the cells that are adjacent within the area returned by the previous step, to create larger boxes. It is noteworthy to point out the tradeoff that arises in this step. Indeed, having a big number of boxes corresponds to a higher number of calls to the Flickr API, which directly impacts on the overall system performance. On the other side, a small number of boxes may affect the efficacy of the LPOIs' identification. As an example, let us consider cities like Venice or Florence, which offer a variety of famous touristic attractions. In these scenarios, it is often the case where the search has a bias towards these well known places considering the amount of photos taken. Counteracting this behaviour, while pursuing our goal, a smaller sized box could be able to capture additional relevant LPOIs. Hence, the system must find a good tradeoff between the efficacy brought by a big number of small boxes and the response time of the overall system.

To this end, the system merges each square, with side equal to the maximum allowed deviation, with its adjacent ones, once in the horizontal direction and once in the vertical direction. In this way, we obtain two sets of rectangles, where each rectangle has one side equal to the maximum allowed deviation. From here on, the set with the smaller cardinality is chosen for the analysis of the contained photos.

In Figure 6 we compare the set of boxes calculated by RouteBoxer (Figure 6(a)) against the ones computed by our algorithm (Figure 6(b)). RouteBoxer tends to produce larger boxes which would reduce the efficacy of the system. Our tests have shown that our partitioning scheme is more effective in finding LPOIs for both horizontal, vertical and diagonal paths.

Once the partitioning is completed, the system sorts the boxes to search for photos. For this purpose the search first starts from even boxes and then considers odd ones, allowing us to populate the initial gallery faster.

From a presentation viewpoint, each collected photo is presented to the user using popup balloons. At this point, the end-user can browse the pictures or choose to play a built-in multimedia presentation with a slideshow of the different photos. The user may choose to select a LPOI as a waypoint: in this case the system calculates the new travel path toward the selection. In addition, the system computes a “proposed travel route” containing all the identified LPOIs. We also allow the user to provide us with some feedback regarding the photos contextualized in the LPOIs by marking the photo as “not relevant”. This information is stored in a database and the system will not select that image anymore as soon as the number of these reports exceeds a threshold. We will return on this point in the following section where we discuss the gamification approach.

Photos might not always be sufficient for selecting a waypoint. To this end, *PhotoTrip* complements the search outcome with additional information originating from the Wikipedia geolocalization API. At this step, textual information is retrieved from the former service by exploiting the geo-coordinates of the photos. A similar service is offered by GeoNames [2]; in contrast to Wikipedia API, GeoNames on average returns a larger set of articles, but the service is constrained on the number of requests per day. Our approach is to switch from Wikipedia to GeoNames if no result is returned from the former[4].

PhotoTrip also allows a radial search which returns local attractions within a maximum deviation from a starting point defined by the user.

4 A Gamification Approach to *PhotoTrip*

A crucial component of *PhotoTrip* is the automatic photo filtering procedure which exploits user-chosen tags to select relevant content from Flickr. This process is imperfect given that the tags on Flickr are chosen by the user and thereby susceptible to the user's perception of the context. Figure 7(a) shows a picture example that we do not want to be returned by the system, since the main subject is the group of nuns in the foreground, hiding the intended subject (the famous Milan Cathedral, or *Duomo di Milano*, in Italy). Figure 7(b) shows a correct photo of the monument.

Partially addressing this problem, we provide the end-users of our system with the capability to return us a feedback which helps to further refine the displayed

[4] In the past we relied on WikiLocation, a service no longer available; see [20] for details.

(a)

(b)

Fig. 6: Travel path segmentation by: (a) RouteBoxer, (b) *PhotoTrip*.

content, improving system precision. However, for the system to scale, we cannot rely entirely on individual users' feedbacks and the main reason for this is their selfish behavior [7].

We might employ artificial intelligence techniques, automatically filtering non-relevant pictures prior to their selection from Flickr. As a consequence, users involvement might be kept to a minimum. However, image and context recognition is a difficult problem where computers have not yet surpassed human intelligence. Moreover, applying a filter to each retrieved image also causes efficiency problems. To this end, we decided to rely on gamification as the way ahead to keep users involved and engaged while performing a task which benefits the community as a whole. The main idea behind the approach is to provide users with an incentive to contribute to the feedback loop, helping the system to select photos of relevance.

4.1 Game description

All the photos retrieved from Flickr are disputed photos which demand feedbacks from users in order to be eligible of being displayed on the search results of the main platform. The goal of the game is to acquire users' feedbacks to improve the results displayed by the system, while providing incentives to the users to engage with the game. Once a disputed photo has reached a majority vote, i.e., a sufficient set of users have voted the photo as relevant or not, the photo becomes eligible of being (or not) displayed whenever a search result encompasses the LPOI it references.

Our choice was to separate the gaming module from the main *PhotoTrip* platform, providing separate interfaces from which the user could switch back and forth. To augment the chances of user involvement and collaboration the game was developed and integrated with Facebook. Indeed, in order to access the gamified application the user is required to login through his/her

(a)

(b)

Fig. 7: Two examples of pictures retrieved from Flickr: (a) a picture that we do not want to be returned by the system; (b) a picture eligible for the system.

Facebook account. The choice of using Facebook as a development platform is motivated by the need of taking advantage of the social and collaborative environments. It is noteworthy to point out that no information related to the user's account or Facebook activity is exploited and all the data gathered relate solely to the game play.

The game is accessed by clicking on the "Play with us" button present in the upper left corner of the platform's landing page. When the user logs in, he/she is presented with some menu choices (leftmost vertical bar) consisting of statistics, achievements, community help, routes and other sections (Figure 8). The most important one is the game play section which in turn is comprised of the following workflow: main landing page, waiting gallery, results gallery and map. When the user chooses to take part in the game, he/she is presented with a search form and with a list of already established challenges in which the user can participate. A new trip search will trigger and schedule a new game play in which the user can engage whereas the established challenges are quests chosen autonomously or manually by a system administrator before being presented to the user. Every time a user searches for a new itinerary, the result is queued in the set of suggested itinerary. Every time the set of pictures corresponding to an itinerary is voted by a sufficient number of users[5], the itinerary is removed from the set of suggestions. Indeed, the challenges might be managed by the system administrator, through his/her interface. We plan to automatically choose the set of suggested itinerary based on the user's vicinity to the challenged routes.

The already established challenges are listed in increasing order based on the number of users that have had an interest in them. The rationale behind this choice is to favor those quests belonging to routes where users have exhibited less interest. Assuming the user triggers a new search, the system goes into the waiting gallery, incrementally presenting the user with the photos related to LPOIs in the search route. When all the results from the search sub-quests are ready, the system goes into the results gallery where the user is presented with a slideshow comprised of a certain number of photos (a system parameter) that require the user's attention. At this point, the user can start voting the relevance of the pictures based upon the photo tags retrieved from Flickr.

The slideshow is comprised of randomly chosen photos from all the LPOIs present in the path. Once a slideshow vote has been completed, the user can either choose to complete the quest for a new slideshow or go to the map displaying all the voting quests he/she has accomplished. As an example, we refer to Figure 9 where the user is presented with a random image selected from the available ones and has to vote it either as being relevant to architecture or not. Once the user casts his/her vote, he/she will be presented with the next photo, if any still remains.

After completing a quest, the user statistics are updated and can be consulted through the appropriate section on the main landing page. The statistics provide an overview related to the user's reputation, ranking, points acquired whose *modus operandi* will be discussed in the following in more detail. Along with these statistics the user can also view how specific photos he/she has evaluated have been voted by the rest of the community. Other crucial components which exploit Facebook to involve others in the feedback loop are the "Help Others" and "Get Help from Others" voices. Indeed, a user can come into the help of other

[5] We discuss in Section 5 how to determine this number.

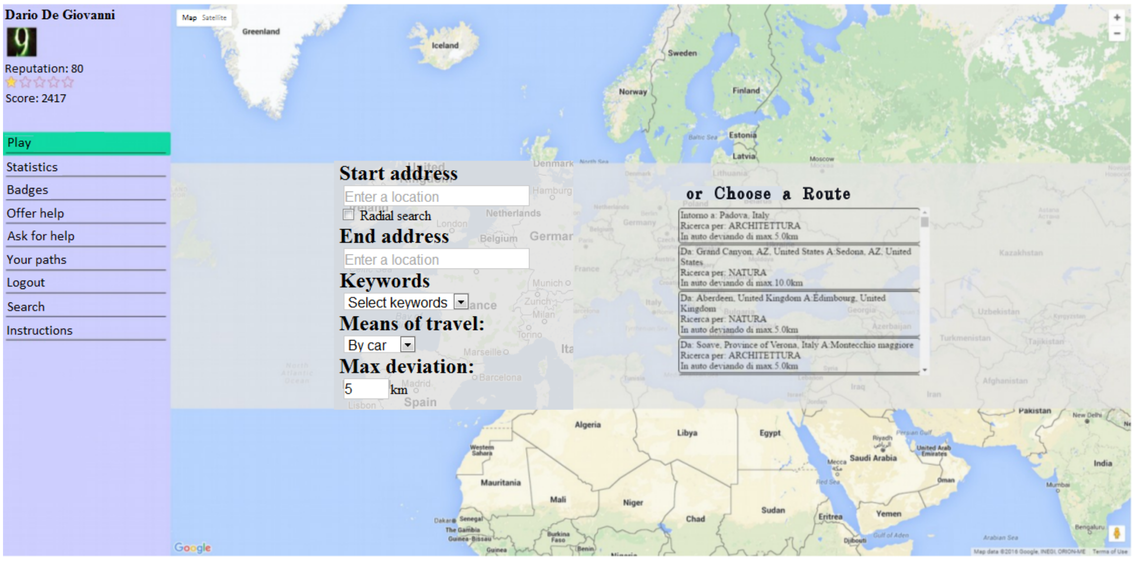

Fig. 8: The gamified version of *PhotoTrip*, with the menu shown in the leftmost vertical bar.

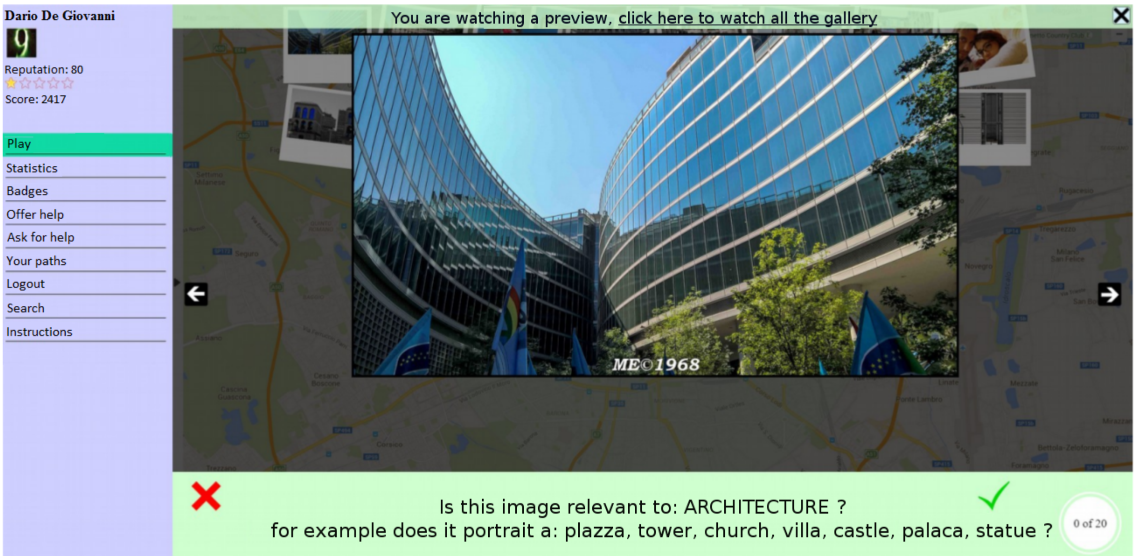

Fig. 9: The user has to follow his/her judgment to vote on whether the image is relevant to "architecture".

users by participating in a voting quest for routes containing disputed photos. These photos might be either fresh photos with no vote attributed or photos without a clear a majority vote. Complementary to the "Help Others" option is the "Get Help from Others" option; if used, it triggers a notification for that specific quest to the user's friend list.

4.2 Majority wins

In the gamified version of *PhotoTrip* we adopt three most common gamification strategies - Points, Badges, and Leaderboards (PBL). Points are provided as scores, corresponding to the measure of each quest the user is involved in and badges are awarded whenever certain thresholds are met. In addition to the first badge which is attributed during the first login of the user, additional ones are attributed when the following criteria are met:

1. a user starts a search for a new path;
2. a user starts m path searches, with m varying among 10, 50 and 100;
3. a user completes a voting quest comprised of n photos, with n=20 used as the default value;
4. a user completes all the voting quests for a path with disputed photos;
5. a user completes a quest for each of the m paths.

Each badge contains a visual icon, a name and unlocking conditions. Participants receive notification when they successfully unlock the badges. The Leaderboard ranks the accumulated performance scores of all participants. In specifics, two rankings are provided: (i) an overall ranking involving all the participating users of the community and (ii) a personal ranking restricted to the user's Facebook friends. In both the rankings each user has the associated reputation level, the username and the accumulated points.

The system uses the majority vote to control the flow of photos from a disputed to an undisputed status which translates to a photo being visible or not in the search results. While participating in a voting quest, the user is not aware of how the others are voting and relies only on his/her opinion. During a quest, a user casts a positive or a negative vote for each presented photo and a majority vote is established by considering the 50% + 1 of the votes. To be precise, as we will discuss later on, the majority vote depends on the user reputation. We must note here that, even if the majority vote can change from "relevant" to "not relevant", users never lose gained points, which are calculated according to the majority vote at the time of voting.

In the beginning, when all the users have the same reputation level, it is possible to reach a majority vote after only an odd number of votes have been casted. Every time a user casts a vote he/she is attributed a number of points as follows:

1. rewarded with 5 points, if the user is the first to vote that particular photo, providing an incentive for the users to get involved in new search quests and introduce new photos into our database;
2. rewarded with 2 points, if the user casts a vote and the vote is on the majority side;
3. rewarded with 0 points, if the user's vote is on the minority side.

When no majority vote exists, it is not possible to attribute the reward to the participating users and the photo is marked as a disputed photo. When a majority vote for that particular photo is reached, the photo is marked as undisputed and the points are rewarded accordingly to all participating users. In the scenario where the photo is a fresh one and the user is the first one to cast his/her vote, the reward of 5 points is attributed immediately.

Different users' votes might not have the same weight; the system gives more weight to a vote casted by a user whose quests end up on the majority side. Intuitively, the majority reward contributes to the users reputation inside the system. This mechanism was inspired by the *StackOverflow*[6] service which follows a PBL approach, stimulating users to contribute. Indeed, the addition of the reputation mechanism tends to provide users with additional stimuli to participate and collaborate in the platform and to discourage users to vote in a wrong manner. The underlying assumption is that the majority vote reflects the true vote, consequently users on the majority side are rewarded the most, and malicioius users who act against the system are ignored as soon as their reputation goes to 0.

A user's reputation is updated each time a user casts a vote; it changes whenever there is a variation of the majority/minority vote for a photo. In this context, a user is rewarded 1 reputation point whenever his/her position for a casted vote changes from minority to draw or from draw to majority. Hence, a user is attributed 2 reputation points if his/her position changes from draw to majority. Conversely, a user's reputation is decreased by 1 when his/her position changes from majority to draw or from draw to minority. In general the reputation points is equal to the number of majority votes minus the number of minority ones, but to avoid recalculating all users' reputation too many times, we also remove two reputation points to users who change a casted vote.

As stated earlier, a user's reputation influences the majority position. The higher the reputation, the higher the influence in the voting process. This is achieved by counting a user's vote as his/her reputation level. The system contemplates for 5 reputation levels which are represented in the user interface through stars. In the beginning, a user is attributed one star and the reputation level changes every 100 majority positions votes. Therefore, a user reaches the 5th reputation level only after 500 majority votes. Considering the penalties discussed above, a user's reputation might even reach level 0 (0 or negative points). In this case, users are attributed no stars.

5 User Evaluation

The system has been online for a period of 4 months overall. It went under two different rounds of tests. During the first one, we wanted to test if the users like/dislike the system to retrieve tourist paths. We collected data during a period of 30 days, during which 110 users have used the service. They were located mainly in Italy, but there has been some access also from the United States, France and England. The average time spent by each user on *PhotoTrip* was about 6 minutes. The system calculated 880 different travel paths, i.e.,

[6] http://stackoverflow.com/

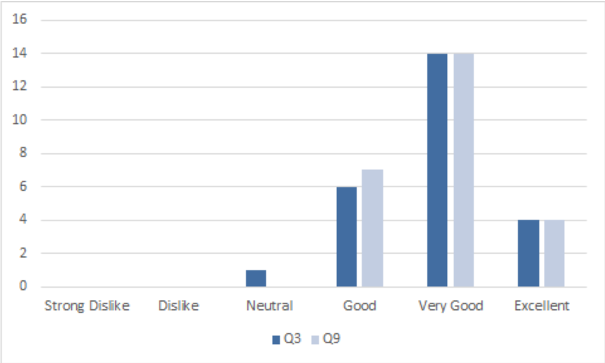

Fig. 10: Answer to questions Q3: "How do you rate the usability of the system?" and Q9: "How do you rate *PhotoTrip*?".

First Test		
Duration	**Users**	**Results**
30 days	110	6 mins per users on average 8 travel paths per users on average
Second Test		
Duration	**Users**	**Results**
10 days	67	38 mins per users on average

Table 1: Summary of the results obtained during the two tests.

on average, each user requested 8 different routes (see Table 1). These data demonstrate the interest raised by our tool; the average time of permanence and the high number of requested travel paths per user show that all users have freely interacted with the site for a quite long time. No problem has been reported with any browser.

We asked the users to answer to a questionnaire of nine multiple choice questions about their experience with *PhotoTrip* [20]. The outcome of the questionnaires has been positive since the great majority of the users reported no difficulty in using the system and considered it very or extremely intuitive (see Q3 in Figure 10). The users were generally satisfied with the output provided by the system. In particular, 72% of the users rated the usability of the system as "very good" or "excellent" and only 4% reported some difficulties. The overall system has been positively judged by 100% of the users with 72% of them describing *PhotoTrip* as "very good" or "excellent" (see Q9 in Figure 10).

We then performed a second round of tests to verify whether the users like/dislike the gamified version of *PhotoTrip*, but also to collect data in order to determine an appropriate threshold value for the minimum number of positive votes needed to show the picture to the users in the provided travel paths. To this end, we asked the users to vote photos from several suggested travel paths, containing both relevant and non-relevant pictures. The considered travel paths varied in terms of search radius and locations. Achieved results are all coherent; for the sake of conciseness, we have hence selected just one representative example to be thoroughly discussed here.

In particular, the representative path discussed here refers to a radial search of 5 km in the surroundings of Padua, Italy, searching for historical or interesting buildings, churches, squares, sculptures, etc.

As anticipated in Section 4, the set of non-relevant pictures is not always easy to identify; we have hence studied the behaviour of the users, in particular in the case of pictures that are not easy to be classified. For instance, Figure 7(a) shows an interesting historical building, but with a group of people in its foreground, thus making it not suitable for our system.

In total, 67 users took part in our second experiment. We divided them into two groups: (i) the "Non-expert group", composed by 34 users who received only basic instructions about how to use the system; (ii) the "expert group", composed by 33 users who also received instructions explaining our interpretation of "relevant image", i.e., pictures that do not contain subjects unrelated to the search tags (e.g., people in the foreground when looking for architecture). Users were free to use the system for as much (or as little) they wanted.

We collected data for 10 days in May 2016, during which the users globally used the system for about 45.2 hours; 38 minutes per user on average (see Table 1). Users tested the system at home, with no one putting pressure on them or making them feel they had to be online for a certain amount of time; the significant time that each user willingly dedicated to the experiment demonstrates the users' interest in the gamified version of the system.

Data collected about how the users voted let us classify the photos into two different sets: those for which users reached a large consensus (a majority quickly emerges) and those collecting contradictory votes (a majority cannot be clearly identified even after many votes). Regarding this aspect, Figure 11 shows how the users voted the pictures for the considered travel path. In particular, the path was characterized by 20 different pictures and we randomly proposed them to the users, one after the other, to be evaluated with respect to their relevance to architecture. Values on the y-axis represent the percentage of users agreeing on a vote (either relevant or non-relevant to architecture), whereas the x-axis represents the considered number of received votes. Thereby, the lines represent the evolution of the percentage of voters that agrees on a vote as the number of voters increases. Values around 50% imply a balance

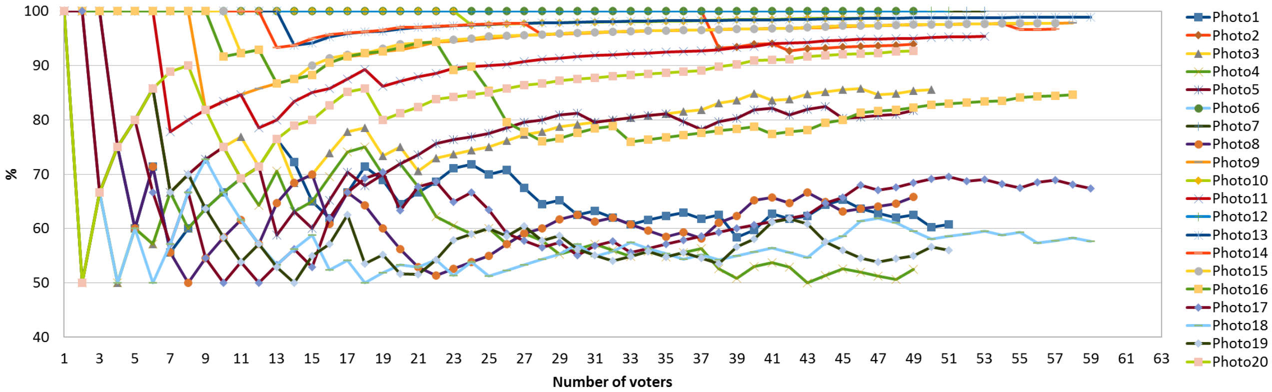

Fig. 11: Evolution of the percentage of voters that agrees on a vote (either relevant or non-relevant) as the number of voters increases; data collected for the travel path from Padua, Italy, for a radial search of 5 km, with keyword "architecture".

between those who cast a positive vote and those who cast a negative one (hence a lack of a clear majority). Instead, values close to 100% correspond to a picture that users can clearly identify as relevant or not to architecture.

Since pictures have been randomly proposed and users could decide how many they wanted to rate, a different number of (different) pictures has been evaluated by a different number of users; this is why the 20 series of values shown in the chart, one for each of the 20 photos, have different lengths over the x-axis.

As we can clearly see from the chart, after a certain (small) number of votes, a set of lines (i.e., photos) is already in the top part of the chart, close to 100%, whereas a second set lays at the bottom of the chart, close to 50%. Figure 12 corresponds to Photo10 plotted in Figure 11 and shows a picture which reached a large majority in a very few votes, as there is nothing related to architecture in the picture. Instead, Figure 13 corresponds to Photo19 plotted in Figure 11 and is a picture that made the users uncertain; in fact, it depicts a sculpture, so it could be considered relevant to keywords about architecture, but the presence of a musician in the foreground makes it also non-relevant for the system.

If we look in details at the data collected about each photo, we can state that for the first set there is not a significant difference in how the "expert" and "non-expert" users voted. For example, for Photo10, 100% of "non-expert" users defined it as non-relevant, and so did 97% of "expert" users. Therefore, both the groups gave the same response. This means that the system does not need "expert" users to find out relevant photos.

Fig. 12: A picture that reached a large majority in a very few votes; the picture was voted as non-relevant to "architecture".

Fig. 13: A picture that made the users uncertain about its relevance to "architecture".

Photo19 was judged relevant by 55.9% of users, but "non-expert" users rated it as relevant (76%) while "expert" users rated it as non-relevant (70%). This means that photos for which the users do not agree with a large majority (more than 70%) should be removed from the system.

These tests showed that common users are effective in helping the system to remove non-relevant photos, even when they have not been previously taught about the meanings of relevant and non-relevant. For these reasons we have included in the system the instructions provided to the "expert" users, but we have decided not to force or stress the users to read them.

Another important result of this test is related to the threshold identifying a majority, i.e., a number of votes considered as sufficient to determine whether a photo is relevant or not. This number is really important since the set of paths is potentially infinite while, on the other hand, we need to convey votes on a limited set of paths (i.e., photos) to quickly decide about their relevance. It is hence crucial to be able to reliably determine when a path should be removed from the set of suggested paths and a picture from the set of photos that can be voted.

To this aim, we decided to set this threshold at 25 votes: if the considered picture reaches a majority of at least 70% relevant votes it remains in the system, otherwise it is removed. We have chosen 25 votes as this value has been sufficient in all the tests we have run to see a clear divergence among relevant/non-relevant pictures and undecided ones (as can be seen even in Figure 11).

6 Conclusion

PhotoTrip is an interactive online service able to autonomously discover and propose to users charming, even if not widely known, cultural heritage locations along travel itineraries. These locations are identified through the analysis of photos contained in the Flickr social network. As Flickr stores more than 350 million geo-tagged photos, *PhotoTrip* can generate travel alternatives anywhere in the world. For instance, Figure 14 shows the results for a radial search within 5 km form Hohenschwangau, Germany, in the case a user asks for LPOIs in the architecture category. As shown, two castles are retrieved: Hohenschwangau Castle and Neuschwanstein Castle, Schwangau, Germany, although German is not in the set of languages supported by the system.

Our system also provides users with additional information on the content of displayed pictures, thus quenching the thirst for knowledge of curious visitors. To this end, Figure 15 shows the information retrieved from Wikipedia about the Neuschwanstein Castle.

The ability to provide for any itinerary, only the most relevant pictures among the many available on Flickr is a crucial component of *PhotoTrip*. Unfortunately, the match between tags and photos on Flickr is strongly affected by the users' context perception. Therefore, we have addressed this limitation by both letting end-users refine the displayed content through feedbacks and using gamification to involve users in the process of improving the response quality of our system. The latter is a particularly interesting approach and we have provided evaluation evidences to demonstrate its efficacy.

To test its compatibility, *PhotoTrip* has been (successfully) combined with different browsers, Google Chrome 18 and later, Mozilla Firefox 9 and later, Internet Explorer 9 and later and Safari 5 and later. Furthermore, the system is accessible through different devices, i.e., desktop, portable PC and tablet. Yet, we plan to improve the interface for smartphones since, due to the small size of the screen, some interaction widgets become too small.

Although Flickr is an international community and its use allows us to search for LPOIs anywhere in the world, we have to report that results tend to be influenced by the language used to specify the tags and the queries. Indeed, photographers usually utilize their mother tongue to describe their photos; only later, sometimes, they add some translations, mostly in English. This also implies that there is no guarantee that tags associated to pictures taken in Italy be in Italian. For this reason, at the moment, we implement two sets of predefined categories, one using the English language and a localized set in Italian.

We plan to extend this work by investigating the possibility to increase the number of retrieved photos through the integration of an online translation service, able to autonomously translate the tags between English and the language of the visited country. This is a crucial issue as it affects even the efficacy of the information retrieval system exploiting Wikipedia or GeoNames, as the majority of the articles present in these web services are in English.

Acknowledgments

The authors would like to thank Michele Torresin and Dario De Giovanni for their technical help during the development of the system and those who took part to the user tests. This work has been partially funded

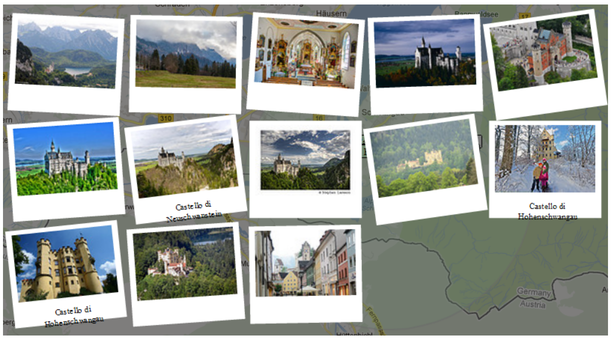

(a)

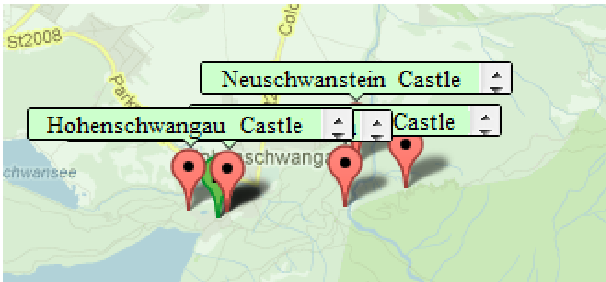

(b)

Fig. 14: Gallery (a) and LPOI (b) returned for a radial search for "architecture" around Hohenschwangau, Germany.

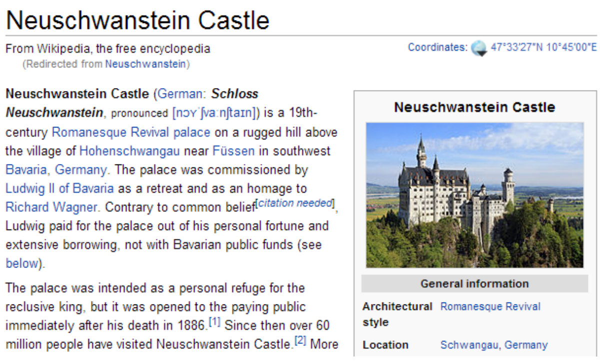

Neuschwanstein Castle

From Wikipedia, the free encyclopedia
(Redirected from Neuschwanstein)

Coordinates: 47°33′27″N 10°45′00″E

Neuschwanstein Castle (German: ***Schloss Neuschwanstein***, pronounced [nɔʏˈʃvaːnʃtaɪn]) is a 19th-century Romanesque Revival palace on a rugged hill above the village of Hohenschwangau near Füssen in southwest Bavaria, Germany. The palace was commissioned by Ludwig II of Bavaria as a retreat and as an homage to Richard Wagner. Contrary to common belief[citation needed], Ludwig paid for the palace out of his personal fortune and extensive borrowing, not with Bavarian public funds (see below).

The palace was intended as a personal refuge for the reclusive king, but it was opened to the paying public immediately after his death in 1886.[1] Since then over 60 million people have visited Neuschwanstein Castle.[2] More

Neuschwanstein Castle

General information

Architectural style	Romanesque Revival
Location	Schwangau, Germany

Fig. 15: Information provided by Wikipedia for Neuschwanstein Castle, Schwangau, Germany.

by the Università degli Studi di Padova, through the projects PRAT CPDA137314 and PRAT CPDA151221.

References

1. Gamification wiki. https://badgeville.com/wiki/Gamification. Last Accessed: July 2016
2. Geonames. http://www.geonames.org. Last Accessed: January 2015
3. Aiolli, F., Ciman, M., Donini, M., Gaggi, O.: Climbtheworld: real-time stairstep counting to increase physical activity. In: Proc. of the International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services, London, United Kingdom, December 2-5, 2014, pp. 218–227 (2014)
4. Amato, F., Chianese, A., Mazzeo, A., Moscato, V., Picariello, A., Piccialli, F.: The Talking Museum Project. IEEE Transactions on Systems, Man, and Cybernetics: Systems (21), 114–121 (2013)
5. Arase, Y., Xie, X., Hara, T., Nishio, S.: Mining People's Trips from Large Scale Geo-tagged Photos. In: Proc. of the International Conference on Multimedia (MM'10), pp. 133–142. ACM, New York, NY, USA (2010)
6. Bellotti, F., Berta, R., De Gloria, A., D'ursi, A., Fiore, V.: A serious game model for cultural heritage. J. Comput. Cult. Herit. pp. 17:1–17:27 (2013)
7. Boyle, E.A., Connolly, T.M., Hainey, T., Boyle, J.M.: Engagement in digital entertainment games: A systematic review. Comput. Hum. Behav. **28**, 771–780 (2012)
8. Chianese, A., Marulli, F., Moscato, V., Piccialli, F.: SmARTweet: A Location-based Smart Application for Exhibits and Museums. In: Proc. of the International Conference on Signal-Image Technology and Internet-Based Systems (SITIS'13), pp. 408–415 (2013)
9. Chianese, A., Marulli, F., Piccialli, F., Benedusi, P., Jung, J.: An Associative Engines Based Approach Supporting Collaborative Analytics in the Internet of Cultural Things. Future Generation Computer Systems (2016). DOI 10.1016/j.future.2016.04.015
10. Chianese, A., Piccialli, F.: Designing a Smart Museum: When Cultural Heritage Joins IoT. In: Proc. of the International Conference on Next Generation Mobile Applications, Services and Technologies (NGMAST'14), pp. 300–306 (2014)
11. Chianese, A., Piccialli, F.: Improving User Experience of Cultural Environment through IoT: The Beauty or the Truth Case Study. Smart Innovation, Systems and Technologies **40**, 11–20 (2015)
12. Coenen, T., Mostmans, L., Naessens, K.: Museus: Case study of a pervasive cultural heritage serious game. J. Comput. Cult. Herit. **6**(2), 8:1–8:19 (2013)
13. Crandall, D.J., Backstrom, L., Huttenlocher, D., Kleinberg, J.: Mapping the World's Photos. In: Proc. of the International Conference on World Wide Web (WWW'09), pp. 761–770. ACM, New York, NY, USA (2009)
14. De Choudhury, M., Feldman, M., Amer-Yahia, S., Golbandi, N., Lempel, R., Yu, C.: Automatic Construction of Travel Itineraries Using Social Breadcrumbs. In: Proc. of the ACM Conference on Hypertext and Hypermedia (HT'10), pp. 35–44. ACM, New York, NY, USA (2010)
15. Deponti, D., Maggiorini, D., Palazzi, C.E.: DroidGlove: An Android-based Application for Wrist Rehabilitation. In: Proc. of the 2009 International Conference on Ultra Modern Telecommunications and Workshops (2009)
16. Deterding, S., Dixon, D., Khaled, R., Nacke, L.: From game design elements to gamefulness: Defining "gamification". In: Proc. of the International Academic MindTrek Conference: Envisioning Future Media Environments, MindTrek '11, pp. 9–15 (2011)
17. Deterding, S., Sicart, M., Nacke, L., O'Hara, K., Dixon, D.: Gamification. using game-design elements in non-gaming contexts. In: Proc. of the CHI '11 Extended Abstracts on Human Factors in Computing Systems, CHI EA '11, pp. 2425–2428 (2011)
18. Froschauer, J., Merkl, D., Arends, M., Goldfarb, D.: Art history concepts at play with thiatro. J. Comput. Cult. Herit. pp. 7:1–7:15 (2013)
19. Furini, M.: On Gamifying the Transcription of Digital Video Lectures. ACM Entertainment Computing **14**, 23–31 (2016)

20. Gaggi, O.: Discovering local attractions from geo-tagged photos. In: Proc. of the Annual ACM Symposium on Applied Computing, SAC '13, pp. 730–735. ACM, New York, NY, USA (2013)
21. Gaggi, O., Ciman, M.: The use of games to help children eyes testing. Multimedia Tools and Applications **75**(6), 3453–3478 (2016). DOI 10.1007/s11042-014-2444-x. URL `http://dx.doi.org/10.1007/s11042-014-2444-x`
22. Google, Inc.: Routeboxer. http://google-maps-utility-library-v3.googlecode.com/svn/trunk/routeboxer/docs/examples.html. Last Accessed: January 2012
23. Hirota, M., Endo, M., Yokoyama, S., Ishikawa, H.: Visualizing shooting spots using geo-tagged photographs from social media sites. In: Proc. of the 15èmes Journées Francophones Extraction et Gestion des Connaissances, EGC 2015, 27-30 Janvier 2015, Luxembourg, pp. 167–178 (2015)
24. Hsieh, H.P., Li, C.T., Lin, S.D.: Triprec: Recommending trip routes from large scale check-in data. In: Proc. of the International Conference on World Wide Web, WWW '12 Companion, pp. 529–530 (2012)
25. Kim, H., An, S., Keum, S., Woo, W.: H-Treasure Hunt: A Location and Object-Based Serious Game for Cultural Heritage Learning at a Historic Site, pp. 561–572 (2015)
26. Kofler, C., Caballero, L., Menendez, M., Occhialini, V., Larson, M.: Near2me: An Authentic and Personalized Social Media-based Recommender for Travel Destinations. In: Proc. of the ACM SIGMM International Workshop on Social Media (WSM'11), pp. 47–52. ACM, New York, NY, USA (2011)
27. Kurashima, T., Iwata, T., Irie, G., Fujimura, K.: Travel Route Recommendation Using Geotags in Photo Sharing Sites. In: Proc. of the ACM International Conference on Information and Knowledge Management (CIKM'10), pp. 579–588. ACM, New York, NY, USA (2010)
28. La Guardia, D., Arrigo, M., Di Giuseppe, O.: A location-based serious game to learn about the culture
29. Lu, X., Wang, C., Yang, J.M., Pang, Y., Zhang, L.: Photo2trip: Generating Travel Routes from Geo-tagged Photos for Trip Planning. In: Proc. of the International Conference on Multimedia (MM'10), pp. 143–152. ACM, New York, NY, USA (2010)
30. Marfia, G., Roccetti, M., Marcomini, A., Bertuccioli, C., Matteucci, G.: Reframing Haute Couture Handcraftship: How to Preserve Artisans' Abilities with Gesture Recognition. In: Proc. of the 9th International Conference on Advances in Computer Entertainment, ACE 2012, pp. 437–444 (2012)
31. Mekler, E.D., Brühlmann, F., Opwis, K., Tuch, A.N.: Disassembling gamification: The effects of points and meaning on user motivation and performance. In: Proc. of the CHI '13 Extended Abstracts on Human Factors in Computing Systems, CHI EA '13, pp. 1137–1142 (2013)
32. Michael, D.R., Chen, S.L.: Serious Games: Games That Educate, Train, and Inform. Muska & Lipman/Premier-Trade (2005)
33. Mirri, S., Prandi, C., Roccetti, M., Salomoni, P.: Food and Gastronomic Heritage: Telling a Story of Eyes and Hands. In: Proc. of the 2016 IEEE Symposium on Computers and Communication, ISCC 2016, pp. 6–9 (2016)
34. Mirri, S., Prandi, C., Salomoni, P.: mPASS: Integrating People Sensing and Crowdsourcing to Map Urban Accessibility. In: Proc. of the IEEE Communications and Networking Conference (CCNC'14), pp. 402–407 (2014)
35. Mirri, S., Salomoni, P., Pizzinelli, A., Roccetti, M., Palazzi, C.E.: "di Piazza in Piazza": Reimagining Cultural Specific Interactions for People-Centered Exhibitions. In: Proc. of the 2016 International Conference on Computing, Networking and Communications (ICNC), pp. 1–5 (2016)
36. Palazzi, C.E., Roccetti, M., Marfia, G.: Realizing the Unexploited Potential of Games on Serious Challenges. ACM Computers in Entertainment **8**, 1–3 (2010)
37. Popescu, A., Grefenstette, G.: Mining Social Media to Create Personalized Recommendations for Tourist Visits. In: Proc. of the International Conference on Computing for Geospatial Research & Applications (COM.Geo'11), pp. 37:1–37:6. ACM, New York, NY, USA (2011)
38. Prandi, C., Ferretti, S., Mirri, S., Salomoni, P.: Trustworthiness in Crowd-Sensed and Sourced Georeferenced Data. In: Proc. of the IEEE International Conference on Pervasive Computing and Communication Workshops (PerCom Workshops), pp. 1261–1266 (2015)
39. Ratan, R., Ritterfeld, U.: Classifying serious games. Serious games: Mechanisms and effects pp. 10–24 (2009)
40. Rego, P., Moreira, P.M., Reis, L.P.: Serious games for rehabilitation: A survey and a classification towards a taxonomy. In: Proc. of the Iberian Conference on Information Systems and Technologies, pp. 1–6 (2010)
41. Roccetti, M., Marfia, G., Bertuccioli, C.: Day and Night at the Museum: Intangible Computer Interfaces for Public Exhibitions. Springer Multimedia Tools and Applications (69), 1131–1157 (2014)
42. Roccetti, M., Marfia, G., Palazzi, C.E.: Entertainment beyond divertissment: using computer games for city road accessibility. ACM Computers in Entertainment **9**, 1–10 (2011)
43. Rosner, D., Roccetti, M., Marfia, G.: The Digitization of Cultural Practices, year=2014, volume=57, issue = 6, pages=82-87. Communications of the ACM
44. Shen, J., Shen, J., Mei, T., Gao, X.: Landmark reranking for smart travel guide systems by combining and analyzing diverse media. IEEE Transactions on Systems, Man, and Cybernetics: Systems (99), 1–13 (2016)
45. Simões, B., De Amicis, R.: Gamification as a Key Enabling Technology for Image Sensing and Content Tagging, pp. 503–513. Springer International Publishing (2016)
46. Sliney, A., Murphy, D.: Jdoc: A serious game for medical learning. In: Proc. of the International Conference on Advances in Computer-Human Interaction, pp. 131–136 (2008)
47. Ueyama, Y., Tamai, M., Arakawa, Y., Yasumoto, K.: Gamification-based incentive mechanism for participatory sensing. In: Proc. of the IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops), pp. 98–103 (2014)
48. Wikipedia: Wikipedia API, API for Wikipedia. https://www.mediawiki.org/wiki/Extension:GeoData#API/. Last Accessed: July 2016
49. Yamasaki, T., Gallagher, A., Chen, T.: Personalized intra- and inter-city travel recommendation using large-scale geotags. In: Proc. of the ACM International Workshop on Geotagging and Its Applications in Multimedia, GeoMM '13, pp. 25–30 (2013)
50. Zheng, Y.T., Yan, S., Zha, Z.J., Li, Y., Zhou, X., Chua, T.S., Jain, R.: Gpsview: A scenic driving route planner. ACM Trans. Multimedia Comput. Commun. Appl. **9**(1), 3:1–3:18 (2013)