

Context-Aware Recommendations Using Mobile P2P

María del Carmen Rodríguez-Hernández
Dept. Computer Science and Systems Engineering
University of Zaragoza
Zaragoza, Spain
692383@unizar.es

Raquel Trillo
Dept. Computer Science and Systems Engineering
University of Zaragoza, I3A
Zaragoza, Spain
raqueltl@unizar.es

Sergio Ilarri
Dept. Computer Science and Systems Engineering
University of Zaragoza, I3A
Zaragoza, Spain
silarri@unizar.es

Ramón Hermoso
Dept. Computer Science and Systems Engineering
University of Zaragoza
Zaragoza, Spain
rhermoso@unizar.es

ABSTRACT

In recent years, there has been an increasing research attention towards Context-Aware Recommender Systems (CARS) for mobile users. The main motivation is that, by considering the current context of the mobile user, more relevant suggestions can be provided. However, further research is required to enable the effective deployment of mobile CARS.

In this paper, we study the possibility to implement mobile CARS by using a pure mobile P2P (Peer-to-Peer) approach, where no centralized database or server exists. Instead, the mobile devices of the users propagate rating information in an opportunistic way, when they become neighbors from a communication point of view. We study the problem by considering a specific use case scenario: the recommendation of items to observe in a museum. Besides, we exploit a synthetic generator of datasets for the evaluation of CARS, called DataGenCARS, to build the recommendation scenarios based on real and synthetic data.

CCS CONCEPTS

• **Information systems** → **Collaborative filtering; Recommendation systems; Mobile information processing systems;**

KEYWORDS

Mobile recommendations, mobile P2P, context-aware recommendation systems

ACM Reference Format:

María del Carmen Rodríguez-Hernández, Sergio Ilarri, Raquel Trillo, and Ramón Hermoso. 2017. Context-Aware Recommendations Using Mobile P2P. In *MoMM '17: The 15th International Conference on Advances in Mobile Computing & Multimedia, December 4–6, 2017, Salzburg, Austria*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3151848.3151856>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MoMM '17, December 4–6, 2017, Salzburg, Austria

© 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-5300-7/17/12...\$15.00

<https://doi.org/10.1145/3151848.3151856>

1 INTRODUCTION

Recommendation systems suggest items that users may find interesting. In this way, they are expected to diminish the cognitive overload that users can experience when they have a myriad of alternatives to choose from. Popular recommendation systems are those provided by Netflix (recommendation of movies and series), Amazon (recommendation of books or other items that can be bought), or Spotify (recommendation of songs), just to cite some popular examples. These systems are interesting both from the perspective of companies and users alike, as long as they are able to offer relevant recommendations to the users.

Moreover, in recent years there has been a great interest in the so-called Context-Aware Recommender Systems (CARS), which have been proposed under the premise that more appropriate recommendations can be offered when the context of the user (e.g., his/her location, time of the day, day of the week, current weather, potential companion, mood status, etc.) is taken into account. These systems add a third dimension (the context) to the traditional 2D recommendation paradigm (users, items). We believe that the context is indeed a key element to consider when offering suggestions to a mobile user. However, the new research avenue on mobile CARS is quite recent and has been relatively unexplored so far.

Specifically, our goal in this paper is to study the potential feasibility of deploying mobile CARS by using exclusively mobile P2P (Peer-to-Peer) communications. Rather than assuming the availability of a centralized server storing a large database of rating information, the idea is that mobile users will propagate partial amounts of rating data in an opportunistic way, that is, when they meet each other in the physical space. Mobile P2P architectures offer several potential advantages over centralized solutions. For example, [21] analyzes some advantages for the specific case of vehicular networks (where the mobile nodes are vehicles), but the advantages can be generalized to other contexts. In the context of recommendation systems, we could highlight, for example, the following: they are costless in terms of the infrastructure required (deploying a fixed support infrastructure may be expensive); the mobile users do not incur any cost derived from the use of cellular communications (e.g., 3G or 4G) when providing rating information; and they may provide better privacy guarantees, as no centralized server collects all the information provided and the transmission

of data is usually confined to nearby spatial regions. So, a P2P approach could provide a number of benefits, as it does not require a fixed support infrastructure that collects all the information, and besides it can always be complemented (if needed) by some fixed support nodes.

As a specific use case, in this paper, we consider the recommendation of items to observe in a museum. For that purpose, we have defined a mixed scenario integrating both real and synthetic data generated by DataGenCARS [12], a synthetic generator of datasets that we have designed for the evaluation of CARS. With a desktop simulator that we have developed, the datasets obtained by using DataGenCARS are processed and ratings are released as needed to simulate votes provided by the visitors to the museum.

The remainder of this paper is structured as follows. In Section 2, we present the proposed approach. In Section 3, we show the experiments performed and the results obtained. In Section 4, we present some related work. Finally, in Section 5, we present our conclusions and outline some prospective lines for future research.

2 APPROACH FOR MOBILE P2P RECOMMENDATIONS

In this section, we present our approach for mobile P2P recommendations. First, we describe the dissemination of data among mobile devices using ad hoc wireless communications. Then, we present the recommendation approach executed locally on each mobile device.

2.1 P2P Data Dissemination

Context-aware recommendation systems exploit a database containing information about items, users, the ratings that the users provide about items, and the existing contexts when the ratings were released. With all this information, they usually try to predict the ratings that the user would provide for items not yet seen and recommend those items whose predicted rating is above a specific recommendation threshold. As mentioned previously, in this paper, we advocate mobile P2P recommendations that exploit exclusively short-range wireless ad hoc communications (e.g., WiFi), which are usually assumed to provide a communication range of 200-300 meters, to exchange data among mobile devices that are within the communication range of each other. In this way, there is no need of a fixed support infrastructure and no centralized server is devoted to collecting all the data (e.g., ratings) provided by the users. Instead, the data are propagated opportunistically (i.e., as the devices meet when they move around the space) through the ad hoc network and the votes provided by the users are stored on the mobile devices of the users in a distributed way. Then, the recommendation system running locally on each mobile device can exploit its local database to predict ratings and recommend items to its user.

More precisely, when a user releases a rating (i.e., provides a vote with his/her assessment of an item), the rating information (user id, item id, context data, and rating provided) is broadcasted in its vicinity: all the mobile devices within its communication range listen to the information and store it in their local databases. Thus, the rating information propagates through the mobile devices of the users according to their spatio-temporal relevance: the rating

information that has been released recently will reach first users located nearby.

The data dissemination approach used to propagate the rating information through the mobile ad hoc network, once it has been released for the first time, is based on *contention-based forwarding*. This implies that, instead of using a flooding approach, which is known to be subject to a number of problems (mainly, the storm broadcasting problem [25], that may cause a major network overload as well as interferences due to network packet collisions), only a single mobile user will be in charge of propagating each specific rating by broadcasting it to all his/her neighbors. All the mobile devices within the communication range listen to the wireless medium and receive the rating information broadcasted, but only one device at a time is responsible for forwarding a specific piece of rating information to other nodes. In this way, the protocol limits the amount of retransmissions of a single piece of data. As suggested by other studies [5, 21], the mobile device of the user located at a furthest distance from the last relay is chosen, in order to maximize the probability that the rating information will travel a large spatial distance as quickly as possible. Indeed, that specific mobile device is expected to have the greatest number of neighbor devices not yet informed about the rating information being transmitted. As an example, Figure 1 illustrates a couple of communication chains.

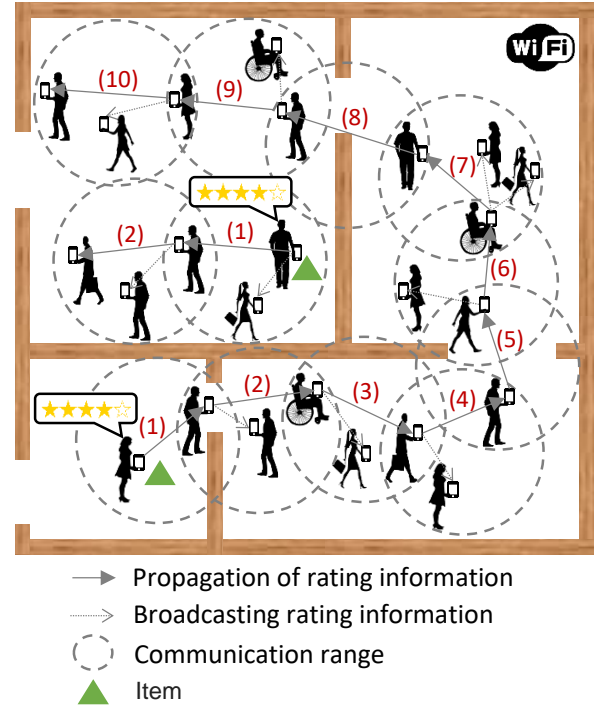


Figure 1: Mobile P2P data propagation.

Moreover, we also allocate a specific *Time-to-Live (TTL)* to each piece of data (e.g., 3 minutes), which controls how long it will be kept alive in the mobile P2P network: when the TTL expires, that rating information is discarded by the current forwarder and not propagated anymore through the network (i.e., it will not be transmitted to other mobile devices). In this way, by adjusting the TTL

value, we can avoid old ratings (released in out-of-date contexts) from further propagation. For each piece of rating information, as long as the TTL is greater than 0, the forwarding responsibility is thus handed over to the furthest mobile device within the communication range. In case there is no neighbor within the communication range, the forwarding role remains assigned to the same mobile device, which will retransmit that piece of data with a certain *retransmission period* until a new data forwarder can be assigned.

The selection of the furthest forwarded is based on the use of *backoff timers* whose duration is inversely proportional to the distance from the last relay. More details about these types of contention-based dissemination protocols can be found in works such as [5, 21].

2.2 Trajectory-Based Recommendation Process

A recommendation application on the mobile device of the user exploits the data stored on the local database of the device, which is filled opportunistically by following the P2P data dissemination mechanism explained in Section 2.1, in order to offer him/her appropriate context-aware recommendations. Figure 2 shows the flow diagram of the proposed trajectory-based recommendation approach. First, through *user-based collaborative filtering (UBCF)*, other users with similar preferences to the user are found and the known ratings provided by those similar users are considered to estimate the potential ratings that the user could provide for different items, in order to determine the top-k items to recommend (i.e., the k items with the best predicted ratings and not yet seen by the user); if the list of candidate items is empty, due to the absence of enough data for the UBCF to provide results (cold start problem), then just the k nearest points of interest are collected as candidate items (this is called the nearest POI, or NPOI strategy). The items with a rating prediction above a *recommendation threshold* (e.g., 2.5, given an evaluation scale from 1 to 5) are added to the list of potential recommendations. Then, the resulting list is re-ordered, if necessary, in order to minimize the distance that the user will need to traverse to access those items. For that purpose, the shortest path passing through all those k items is computed and that path is the one recommended to the user.

Moreover, as illustrated at the bottom of Figure 2, there are several circumstances that could trigger a reevaluation of the recommendation process: 1) when the recommendation system executing on the mobile device has significant new information (it has new data regarding at least a certain number of ratings, provided by the user himself/herself or received from other users, according to the required *knowledge base increase threshold*), 2) when the user is about to leave an area such as a room in a building (and therefore it is convenient to check if there is any additional item in that area worth visiting at that moment), and 3) when the user has deviated from the recommended trajectory significantly (e.g., because something not recommended attracted his/her attention). Besides, in order to avoid recommendation instability, a *minimum time interval between successive recommendation updates* is considered. Finally, independently of the previous conditions, the recommended list of items is updated if the list becomes empty because the user

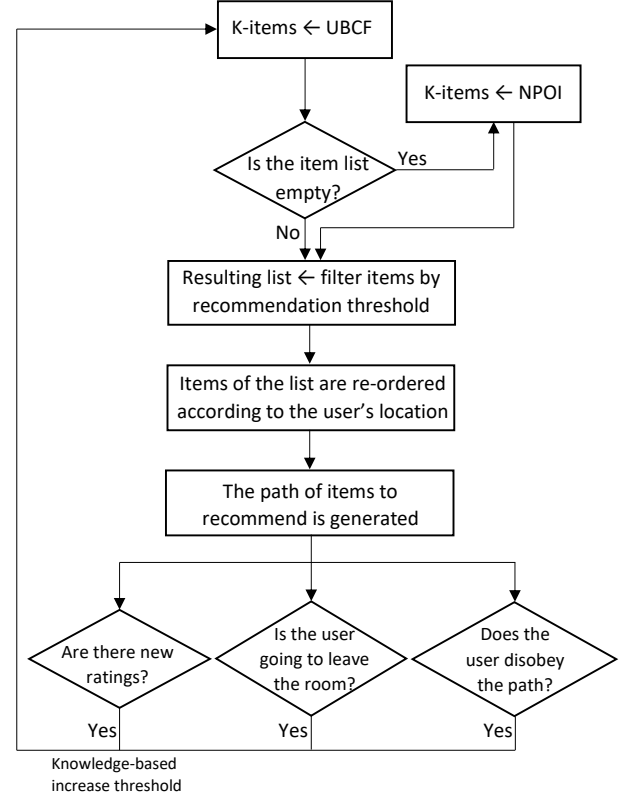


Figure 2: Recommendation process.

already observed all the items previously recommended. The previous rules ensure that an appropriate list of recommendations can be automatically maintained up-to-date in a suitable way.

3 EXPERIMENTAL EVALUATION

In this section, we discuss the experimental evaluation performed to show the feasibility and potential interest of a mobile P2P recommendation approach. First, in Section 3.1, we describe the use case scenario considered. Then, in Section 3.2, we present the experimental settings. Finally, in Section 3.3, we describe the results obtained.

3.1 Use Case Scenario

As a use case scenario, we considered that of user guidance in a museum, which we introduced in a previous seminal work focused on a centralized recommendation approach [13]. The goal of the recommendation system is to maximize the user's satisfaction with the visit. Specifically, the system will suggest to the user a trajectory to follow within the museum, prioritizing the sequence of works of art to observe in the museum within the visiting time available. We were inspired by the Museum of Modern Art (MoMA) in New York and tried to reproduce some of its elements. In particular:

- We used a real dataset composed by information about the works of art exhibited in that museum [26]. The dataset contains 129,024 works of art, with a total of 29 attributes, and

belonging to 14, 949 different artists. The works of art include basic metadata, such as the title of the work, its dimensions, its author and his/her nationality, etc. In order to facilitate the management of this dataset, we performed some pre-processing tasks (e.g., discretization of some attributes, and binarization of others). Unfortunately, the dataset does not contain information about the precise locations of the works of art, which is an important attribute for a context-aware recommendation system. Therefore, we used the Java-based synthetic data generator DataGenCARS [12], that we developed for the evaluation of context-aware recommendation systems, to assign a *geographic location* for each work of art. The locations were generated randomly by complying with specific conditions that ensure a feasible real-world distribution (e.g., they are distributed uniformly along the walls and interiors of the rooms, in order to avoid overcrowded areas or situating two works of art too close to each other). Besides, for each piece of art, we also generated a synthetic attribute representing the *emotion transmitted* by that piece (happiness, sadness, or neutrality).

- We mimicked part of the layout of the MoMA museum. For simplicity, we focused on floors fourth and fifth, which contain a total of 240 items (paintings and sculptures). A desktop simulation application, that we developed by using the tools shown in Table 1, is able to load and display the layout of those floors and their works of art (as an example, see Figure 3).

Tool/library	Use
Java 8	Programming language
Swing	GUI widget toolkit for Java
WebPlotDigitizer [31]	Obtention of key locations (e.g., rooms, doors, and stairs) from raster map images of the museum
JGraphX [6]	Management and visualization of graph structures
Sqlite-jdbc [34]	Management of SQLite databases

Table 1: Tools used for the implementation of the museum simulation application.

- The MoMA dataset contains only information about the works of art in the museum, not about ratings provided by users. Therefore, we used DataGenCARS to generate rating data enriched with context information (i.e., each rating generated is associated to information about the existing context associated to that rating, as each rating is provided when a user observes a work of art in a specific context). We defined several *user profiles*, which are scoring functions that are used exclusively to synthetically generate a rating based on a set of attribute values [12], and assigned a specific user profile to each user simulated. These profiles are obviously unknown to the recommendation system and include context attributes that might affect the ratings provided (see Table 2). For each user, according to his/her profile, a rating was generated for each work of art in the museum, in each possible context.

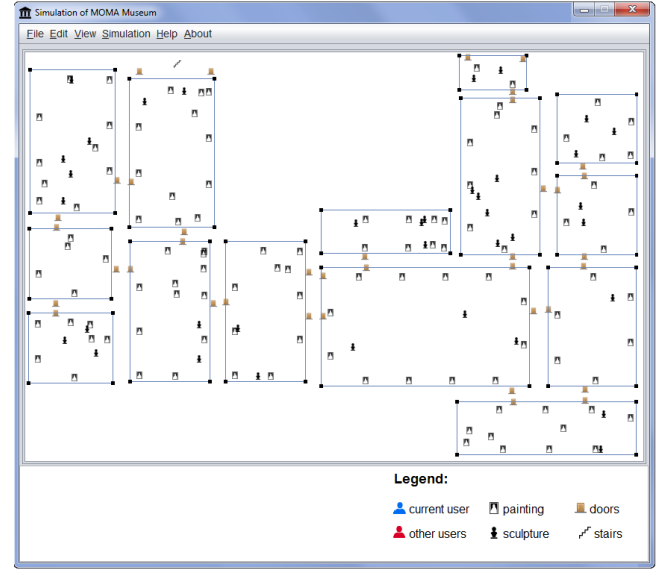


Figure 3: Simulation application showing floor 5.

Context attribute	Possible values
User's mood	Happy, sad, or neutral
Temperature of a room	Warm, hot, or cold
Number of people in a room	Large, medium, or small
Noise level in a room	High, medium, or low

Table 2: Context attributes considered.

- The desktop simulation application developed allows to simulate visitors moving throughout the museum, observing works of art, and releasing ratings. When a visitor's trajectory stops at a certain work of art, he/she observes it for a certain amount of time and then emits a rating according to his/her satisfaction with that work and the current context. As commented before, the ratings were previously generated using DataGenCARS, and the simulator just needs to select the correct one according to the specific context existing at that time. A special case concerns the user's mood, as its impact on the rating provided depends on the emotion transmitted by the work of art observed; therefore, we defined a simple rating adjustment function that can slightly modify the rating generated depending on the relation between the current mood of the user and the emotion transmitted by that work of art.

Figure 4 shows a snapshot of the simulation application at a certain time instant, where we can observe different users observing works of art or moving to other locations. In the figure, there are no doors connecting rooms in the left part of the figure and the right part of the figure: the reason is that the left and the right part of the figure correspond to different floors (connected through the stairs), but for the sake of simplicity we show them next to each other.

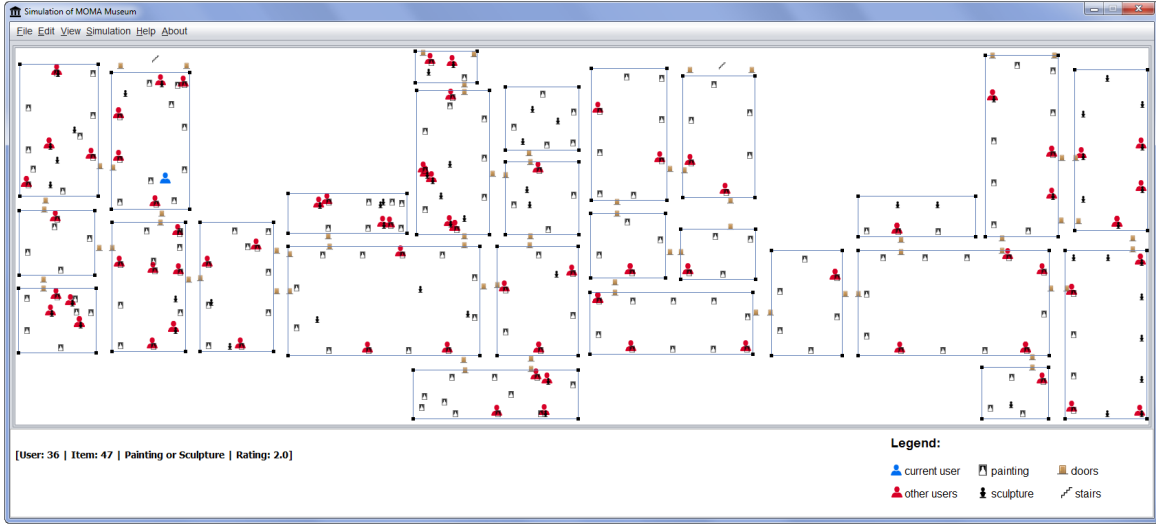


Figure 4: Simulation application in action: visitors in the museum.

3.2 Experimental Settings

The evaluated user is initially located at the front door of the museum and follows a path suggested by the context-aware recommendation application installed on his/her mobile device, while the other museum visitors start from a random location within the museum and follow other synthetically-generated trajectories. The experimental settings used for evaluation are shown in Table 3. In some of the experiments, we compared our proposed P2P context-aware recommendation approach with several recommendation baselines:

- *Completely-random recommender (FULLY-RAND)*: the user visits works of art recommended in a completely random manner, even if this means changing from one room to another that may be located far away. This is expected to be the worst approach possible, as the user could potentially have to traverse very large distances between observations.
- *Exhaustive visit recommender (ALL)*: the user is recommended to visit all the works of art in his/her current room, then the user is suggested to move to a different connected room (in the same or in a different floor, if there are stairs nearby), and so on.
- *Nearest Point Of Interest recommender (NPOI)*: the user is recommended to go to the nearest point of interest, which may be a work of art or an exit (a door connecting to another room, or the stairs leading to a different floor).
- *Centralized recommender (Centralized)* [13]: a central server stores information about all the ratings that are released along time and, based on all these data available, a user-based collaborative filtering strategy is applied. The k items with the highest estimated rating and exceeding the predefined recommendation threshold are re-ordered in order to minimize the distance traversed by the user, and the sorted list is finally recommended to the user.
- *Know-It-All recommender (Know-It-All)*: it is like the centralized approach but assuming that the centralized service

stores all the rating information that the other visitors could eventually provide for each work of art in the museum (even if they actually never released that information or observed those items). This is obviously an unrealistic assumption, and therefore this approach is a baseline operating under ideal conditions.

- *K-Ideal recommender (K-Ideal)*: it is similar to Know-It-All but, rather than applying a user-based collaborative filtering, all the ratings about items unseen by the user are estimated and the k items with the best predicted ratings are recommended, independently of whether they exceed or not the recommendation threshold. The k -items are re-ordered, if needed, to minimize the total distance traversed by the user.

All the strategies were implemented within the context of the generic framework that we proposed in [11] and using Apache Mahout (<http://mahout.apache.org/>). For the strategies based on collaborative filtering, in order to alleviate the cold start problem, when there is not enough data to obtain a recommendation we apply a default NPOI strategy, as described in Section 2.2.

3.3 Experimental Results

In Figure 5, we show the average of the ratings provided by the user for the items he/she observed in the museum, considering a rating scale in the range of 1 to 5. For the mobile P2P approach, we considered two possible cases: one where the TTL of the data transmitted is set to 3 minutes and another one where the TTL is ∞ (i.e., all the data are kept alive in the network during the whole simulation). In both cases, we observe that the P2P strategy achieves a higher average rating than most strategies (FULLY-RAND, NPOI, and ALL). The only exceptions are the *Centralized* strategy, that behaves slightly better than mobile P2P thanks to a higher availability of information (all the data are collected by a centralized server and made available to the recommendation process), and Know-It-All and K-Ideal, which are unrealistic ideal alternatives. Although the differences between the ALL and P2P strategies do

Parameter	Default value
Number of simultaneous visitors (floors 4 and 5) & visiting time in the museum	176 visitors (inspired by [37]) & 1 hour
Number of items & visitor's average speed & observation time (of a painting or sculpture)	240 items & 3 Km/h & 30 seconds
Trajectories followed by other visitors in the museum	50% of the visitors moving always to the nearest POI (NPOI strategy) 50% of the visitors observing works exhaustively (ALL strategy)
Time needed to change to another floor (take the stairs or the elevator)	60 seconds
Number of recommended items to keep in the result list (K)	10 items
Similarity threshold for the UBCF algorithm	0.5 (Pearson correlation)
Recommendation threshold (1-5)	2.5
Knowledge base increase threshold	40 new ratings
Minimum time interval between successive recommendation updates	30 seconds
TTL of the data to propagate	3 minutes
Communication latency	1 second
Communication bandwidth	54 Mbps (IEEE 802.11g)
Communication range	250 meters
Communication obstacles	Walls block signals (each data communication limited to a single room)
Retransmission period	1 second

Table 3: Experimental settings.

not seem very significant, the ALL strategy implies that the user sees more items that he/she does not particularly like (*no likes*), as we will explain below with Figure 6.




Figure 5: Average rating provided for the items observed.

Figure 6 shows the number of *likes* (items observed and particularly well rated by the user with a score no smaller than 3.5 in a 1-5 scale), *no likes* (items rated by the user with a score below 3.5), and the difference between them. The P2P variants exhibit a performance close to the centralized strategy and are only outperformed by that centralized approach and by the ideal and unrealistic strategies. Although NPOI achieves a high number of likes, it also leads to a considerably high number of no likes, and so an overall higher dissatisfaction. Besides, as it was shown in Figure 5, the average rating obtained with NPOI is also lower than with other strategies.




Figure 6: Likes and no likes.

We also analyzed the evolution along time of the proposed mobile P2P recommendation approach. First, in Figure 7 we can see the

number of votes released along time. Figure 8 depicts the MAE (Mean Absolute Error) in the predicted rating of the items observed by the user during his/her visit to the museum. The dashed trend line depicted in Figure 8 shows that, as expected, the quality of the recommendations improves along time (i.e., the MAE decreases), as the knowledge base stored on the user's mobile device (initially empty) increases its size thanks to the mobile P2P communications with other users.




Figure 7: Votes provided by the visitors along time.




Figure 8: Prediction error along time.

Results for other experiments are omitted due to space constraints. For example, we have also evaluated the impact of other

intermediate TTL values (between 3 minutes and ∞) on the performance of the mobile P2P recommendation approach: higher TTL values eventually lead to larger local databases and a higher recommendation accuracy, at the expense of a higher number of network communications.

As a conclusion, the experimental results show that the mobile P2P recommendation approach provides satisfactory results while avoiding the potential inconveniences of a centralized approach. Obviously, the centralized approach is slightly better because it applies the same recommendation process but it exploits a larger database to learn appropriate recommendation models; instead, the mobile P2P alternative relies on a local database with information collected opportunistically by the mobile device, rather than a centralized database with all the information generated until then.

4 RELATED WORK

There exist many approaches on recommendation systems [30], that can be applied in a variety of domains, such as the recommendation of movies, music, Points Of Interest (POIs), etc. These systems try to alleviate the overload of users when they need to choose among a large amount of potential options, by suggesting them items that they could find interesting. In the last years, with the emergence of CARS, the context has been taken as a fundamental piece to allow those systems to build more personalized and efficient recommendation mechanisms, especially in tourism [16]. In this section, we analyze some related work. First, we focus on mobile recommendation systems in general. Then, we consider in particular those related to mobile P2P environments. Finally, we describe some works related to the example use case considered in this paper, which is that of museum guidance.

In this study of related work, we omit simulation approaches, as this is not the focus of our paper. However, it is interesting to mention that an interesting approach on mobility analysis in buildings is presented in [4]. In that work, the authors use a pedestrian simulator in order to study how the mobility of people may affect a building's functioning, with the goal of enhancing different services of the building, such as exit routes in case of an emergency evacuation. Although the simulation approach is related to the work presented in this paper, in our case we use DataGenCARS as a tool that allows to generate data that is suitable for evaluations related to context-aware recommendation scenarios. Furthermore, our approach pursues the enhancement of the individual user's experience, rather than the performance of a facility as a whole.

4.1 Mobile Recommendation Systems

Mobile recommender systems provide recommendations for mobile users [29]. An example of mobile CARS is the one presented in [10]. In this work, the authors put forward a recommendation system for mobile devices based on the use of semantic web technologies in order to provide users with suggestions about the best movies to watch and the most appropriate theaters to go to. To calculate the recommendations, they use context data such as the time, the user's location, or crowd measures. Although the approach takes those factors into consideration, the authors do not deal with the propagation of information among different users, as we do in our work. We claim that this information may help the recommendation system

to provide more accurate recommendations in real time, since that information somehow contains what is currently occurring in the scenario considered.

Regarding the type of environment to deal with, we find different approaches [15, 19, 24]. For instance, in [15], a recommendation system approach for mobile devices is presented focusing on open-air environments, where the users can freely move. The system provides sequences of POIs in the environment that could be of potential interest for the user, as well as recommended tours. The concept of context in this system is merely the use of information about the remaining time available to visit the area. Although this work represents a remarkable advance in the state-of-art of the field, we strongly believe that the context should also be enriched with information provided by other participants. In our case, the users are able to exchange opinions about what they have seen so far, so allowing the recommender system to build more accurate recommendations.

Along the same line, a novel approach for recommending tourist routes can be found in [19]. Its authors aim to provide the best routes for users by taking into account their past activity and the path's overall distance, as well as their psychological preferences in the route. They tackle this issue as a multi-objective optimization problem. However, this work does not take into consideration what we believe are important issues, such as the environmental context or the potential of the dissemination of information among users.

Finally, another work to be highlighted is found in [24], which presents an approach similar to ours, since it intends to build routes for users in a shopping mall by taking into account contextual information of both retailers and visitors. An interesting property of this work is the consideration of the friends' opinions as part of the recommendation process. In our case, we take a significant step further, since we allow users to exchange opinions in a decentralized fashion. Moreover, while friends may have opinions that are out-of-date or biased due to a potentially-reduced circle of friends, with our approach the current visitors' opinions provide a clear picture of what is liked or disliked in real time and taking into account recent context information.

Privacy is another important issue when communicating opinions to calculate recommendations. An example can be found in [28], where the authors propose a method to preserve the privacy of users when communicating their opinions to a centralized recommender system. Mobile devices generate, on behalf of the users, fake opinions that are sent along with the real ones to the recommender. When recommendations are sent back to the mobile device, those considered to match with the fake information previously sent are discarded.

A similar approach is put forward in [38], in which the authors propose a distributed collaborative strategy for the recommendation of location services by using the neighboring users' location information profiles.

Although the development of privacy protection measures is out of the scope of this paper, in order to protect the confidentiality of information exchanged by users there exist some approaches that use encryption mechanisms to keep the confidentiality of information. For example, in [38], the authors propose a distributed collaborative strategy for the recommendation of location services by using the neighboring users' location information profiles. In our

approach, the distributed nature of the mechanism makes it difficult to incorporate traditional methods, that frequently require a central point of control at least at some step of the process. However, as our proposal is based on mobile P2P interactions, sensitive data are never collected by a server in a centralized way: instead, the data are transmitted distributively usually in the spatial proximity.

4.2 Mobile P2P Recommendation Systems

The distributed nature of our approach makes it a relevant case in the literature on recommendation systems. Thus, most recommendation systems are implemented in a centralized way. However, some researchers have also considered the P2P architecture as an interesting alternative to reduce the calculation complexity of centralized recommendation algorithms. One of the first approaches that addressed the problem of scalability in recommendation systems was presented in [36]. The authors proposed a P2P-based collaborative filtering approach that suggests products and services to mobile customers. Recommendation queries are broadcasted to neighbor peers, which are software assistant agents interacting with a mobile customer. In the same application domain, Zhang et al. [41] proposed a distributed recommendation approach based on cloud computing. In this distributed architecture, the authors divided the rating database into sub-databases (by using the K-means clustering method), which were stored in distributed cloud servers. The recommendation queries are submitted to the different cloud servers in a collaborative way and the data are updated by using a Distributed Hash Table (DHT). As opposed to these proposals, we consider each mobile device as a peer that contains rating information generated by the user himself/herself or obtained from other nearby peers (or mobile devices) in an opportunistic way.

Another early example [18], where the authors propose a P2P architecture to achieve a reasonable storage structure for opinions about items. In this proposal, the users maintain information as simple containers; specifically, DHTs are used and new entries location are calculated measuring similarities. However, in our approach users exchange information with other users who are currently in the same area, so those opinions exchanged can vary the resulting recommendations.

Another work worth mentioning is in [27], where the authors present a P2P recommendation system that exploits trust as a cornerstone for users to share their information with others. In this context, the trust between two users is interpreted as their capability to predict the ratings of common choices. As opposed to our proposal in this work, rather than using an opportunistic approach for the propagation of rating information, they rely on flooding queries through the network up to a specific hop distance.

In [23], the authors presented a P2P recommendation approach different from ours. It is based on a naïve Bayesian classifier and uses a P2P topology that preserves the privacy of users. Through a P2P network, the users can communicate with each other and exchange data to generate predictions. For example, when a specific peer needs to predict the rating of an item, it sends a request of ratings about that specific item to other peers in the network. The peers which rated that item predict the probability values for each class (likes or dislikes) and then send the calculated values to the peer submitting the query. Besides, in order to preserve the privacy,

the authors applied a specific mechanism to hide rated items in each peer. This proposal is quite different from ours because it implies explicitly querying for ratings when needed, rather than receiving rating information opportunistically when the mobile peers meet each other. A challenge for the practical application of this approach in a mobile P2P network is how to ensure the routing of the answers to the originating peer by using only wireless ad hoc communications, which is difficult due to the continuous movement of the mobile peers.

With the same flavor, an interesting P2P recommendation system for large-scale data sharing is presented in [14]. This work puts forward a mechanism that is able to communicate information to other peers based on a semantics-based gossip protocol. Moreover, the authors also propose a query routing algorithm able to identify relevant peers, in order to avoid the problem of information flooding, by calculating similarities among peers and topics. In our case, we decided to use content-based forwarding and a TTL field to avoid the potential network overhead.

Another relevant example is the *iTravel* system, that recommends attractions to tourists in a mobile environment [39]. *iTravel* supports mobile P2P interactions, using short-range wireless communications such as Wi-Fi or Bluetooth, to exchange the ratings of the visited attractions among the mobile devices of the users. The authors presented three data exchange methods. In the first method (unconditional data exchange), when the recommendation application of a user detects another *iTravel* application nearby, it sends to that *iTravel* application all the ratings stored in its database; to avoid the information flooding problem, the maximum number of propagations of ratings is limited. The second method (preference-based data exchange) avoids the information flooding problem by considering the user's preferences when deciding which information to exchange: *iTravel* propagates only the rating lists of similar users (i.e., rating lists that are similar to the recipient's rating list). This alternative may be inappropriate for the case where the user is surrounded by other users with tastes different from their own. Hence, the third proposed method (hybrid data exchange) is a combination of the two previous ones. As opposed to this work, we avoid the information flooding problem by propagating the ratings to the furthest mobile device and by using a TTL field. Besides, in our approach the ratings are propagated opportunistically through the network of mobile users when they are released.

A context-aware recommendation system, called *PPNews*, that proactively pushes news in a mobile P2P network, is presented in [40]. The authors implemented it by using *JHPeer*, which is a hybrid P2P framework built by themselves, that supports peer-to-peer communications. *JHPeer* provides basic context-aware services to handle context queries or deliver contextual information to peers. During the recommendation process, first the mobile device establishes communication with one of the peer servers to communicate the user's profile information. Then, the context service (*JHPeer*) installed on the user's mobile device automatically delivers context information (e.g., location and usage pattern) to the server. Finally, the server proactively recommends the top-k news articles to the mobile user based on the user's profile, the context information, the news content, and the ratings stored on the peer. The rating prediction is determined by using a hybrid recommendation approach (content-based filtering and collaborative filtering). As opposed to

our work, this proposal uses a hybrid P2P network, with a set of servers that provide load balancing and support distributed services. These servers gather real-time news articles from providers via RSS.

4.3 Recommendation Systems for Museums

Providing appropriate recommendations for museum visitors has also been a hot topic in the field of recommender systems, and so we can find a significant number of works dealing with it. Some of them focus on the design of systems that serve as the visitors' guide trying to offer a personalized experience. For instance, a couple of collaborative models to predict the upcoming locations of users in a museum have been proposed [3], which could be used to build user models; more specifically, a temporal approach is proposed to predict future locations based on the potential interest of unseen exhibits, and a transitional approach is presented to predict future locations based on the paths followed by other visitors. A collaborative filtering algorithm to recommend trajectories to visit POIs, which is evaluated with data regarding the Vienna Zoo (Austria), has also been presented [20]. As contextual information, this system takes into account the locations of the users, and so the authors present a notion of similarity between contexts. Although this work presents rather interesting features, it has also some drawbacks (e.g., it does not consider the opinions of other visitors, in terms of ratings provided rather than just POIs visited, and its evaluation is quite limited). Also on the use of collaborative filtering to recommend attractions to visitors, we find an appealing approach about POI recommendation taking into account the location diversity [7]. This work is evaluated by using outdoors check-in app-related datasets where the only contextual information provided is the location of the users. In our opinion, those datasets somehow limit the different types of scenarios where recommender systems could be tested on, especially when dealing with CARS, as the context information available should be rich instead. *MusA* [32] is a generic framework to develop multimedia guides for mobile devices and provides a vision-based indoor positioning system, as well as thematic paths created by professional curators or museum staff. Another interesting work is *SmARTweet* [8], a location-based application able to detect visitors' nearby artworks and show their corresponding history using multimedia resources. It also uses collaborative filtering to recommend personalized artworks based on the visitors' interests (gathered via questionnaires) and behavioral information collected by tracking the users throughout the museum.

There exist some other approaches also aimed at building reliable guides for visitors in museums. For example, *UbiCicero* [17] provides recommendations thanks to the detection of artworks nearby by using RFID readers assembled in mobile devices. It also takes into account some contextual information, such as the visitor's current position and his/her behavioral history. With this information, the system returns personalized recommendations. The evaluation of this proposal is quite limited too, since the system is tested with groups of 5-7 users. A recommender system for mobile devices for museum visitors presented in [2] can be seen as a step forward, as it is able to adapt to the user's interests, by taking into account contextual information such as the location, expertise, and time. It uses a hybrid mechanism based on collaborative filtering combined with a post-filtering semantics-based approach.

When focusing on the way these systems are tested, we realize that there is no common methodology. To our knowledge, there does not exist any reliable dataset for testing CARS, especially in indoor environments. Thus, some works deploy real implementations and test the suitability of the systems by asking the visitor to fill out questionnaires when leaving the museum, in order to know their satisfaction [1, 17, 32, 33]. Although this method might be a good way to evaluate this type of systems, the real outcomes seem to be very simplistic, since collecting a significant and reliable set of responses is too costly; moreover, even if the visitor is satisfied, it could be difficult to know if she/he could be even more satisfied with a different set of recommendations. We can also find some approaches that do not test their approach but present different case studies [2, 9, 22, 35].

As opposed to the existing work, our goal in this paper is to show the potential benefits of a context-aware trajectory-based recommendation system, that exploits ratings provided by other visitors in real-time and using ad hoc wireless communications. Moreover, we propose an experimental evaluation method that makes use of real data enriched with synthetic data generated by using the DataGenCARS tool. This approach provides virtually unlimited evaluation opportunities with a low cost in comparison with performing evaluations with real users in a museum, and could be potentially used to evaluate any other recommendation approach proposed in the literature.

5 CONCLUSIONS AND FUTURE WORK

In this paper, we have analyzed the possibility to deploy a context-aware recommendation system using mobile P2P ad hoc communications. The proposal is to use a contention-based forwarding approach to disseminate, in an opportunistic way, the information about ratings when they are released by the users. We have presented experiments with a use-case scenario built from both real data and synthetic data created by using a synthetic data generator for the evaluation of CARS, called DataGenCARS. The results obtained show that it is possible to deploy a relevant trajectory-based recommendation strategy based on collaborative filtering and without the need of centralized servers. Besides, the use of the DataGenCARS to build data that are later exploited by a simulator of the scenario has enabled the evaluation of the proposal without the high cost that deploying and evaluating it in a real context would have implied.

The experiments presented in the paper are a set of representative results. Through other experiments using user trajectories that lead to highly-skewed distributions of visitors in the museum, we have observed that in those cases the cold start problem might re-appear when there is a major change in the user's location: as the information propagated through the ad hoc network tends to spread over the spatial area only gradually, reaching an area with a small number of visitors could be problematic at the beginning; for example, if the user moves to a floor of the museum where the number of visitors is very small, it could happen that the local database on his/her mobile device initially has little information about the works of art in that floor, which would lead to poor predictions about those works. A mixed P2P approach, based on some support fixed nodes located at strategic places and storing some

information about ratings released in nearby areas, could be used to help to diminish this problem. As future work, we could analyze this issue in more detail. Besides, it would be interesting to perform additional tests considering other application scenarios.

ACKNOWLEDGEMENTS

This research has been supported by the projects TIN2013-46238-C4-4-R, TIN2016-78011-C4-3-R (AEI/FEDER, UE), and DGA-FSE.

REFERENCES

- [1] Inmaculada Ayala, Mercedes Amor, Mónica Pinto, Lidia Fuentes, and Nadia Gámez. 2014. iMuseumA: An Agent-Based Context-Aware Intelligent Museum System. *Sensors* 14, 11 (2014), 21213–21246.
- [2] Idir Benouaret and Dominique Lenne. 2015. Personalizing the Museum Experience through Context-Aware Recommendations. In *International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, Kowloon (China), 743–748.
- [3] Fabian Bohnert, Ingrid Zukerman, Shlomo Berkovsky, Timothy Baldwin, and Liz Sonenberg. 2008. Using Interest and Transition Models to Predict Visitor Locations in Museums. *AI Communications* 21, 2–3 (2008), 195–202.
- [4] M. Brambilla and L. Cattelani. 2009. Mobility analysis inside buildings using Distrimobs simulator: A case study. *Building and Environment* 44, 3 (2009), 595–604.
- [5] Nicolas Cenerario, Thierry Delot, and Sergio Ilarri. 2011. A Content-Based Dissemination Protocol for VANETs: Exploiting the Encounter Probability. *IEEE Transactions on Intelligent Transportation Systems* 12, 3 (September 2011), 771–782.
- [6] Chen Chao. 2016. JGraphX. (2016). <https://github.com/jgraph/jgraphx>, [Accessed on August 16, 2017].
- [7] Xuefeng Chen, Yifeng Zeng, Gao Cong, Shengchao Qin, Yanping Xiang, and Yuanshun Dai. 2015. On Information Coverage for Location Category Based Point-of-interest Recommendation. In *29th AAAI Conference on Artificial Intelligence (AAAI)*. AAAI Press, Austin, Texas (USA), 37–43.
- [8] Angelo Chianese, Fiammetta Marulli, Vincenzo Moscato, and Francesco Piccialli. 2013. SmartTweet: A Location-Based Smart Application for Exhibits and Museums. In *International Conference on Signal-Image Technology & Internet-Based Systems (SITIS)*. IEEE, Kyoto (Japan), 408–415.
- [9] Shih-Chun Chou, Wen-Tai Hsieh, Fabien L. Gandon, and Norman M. Sadeh. 2005. Semantic Web Technologies for Context-Aware Museum Tour Guide Applications. In *19th International Conference on Advanced Information Networking and Applications (AINA)*, Vol. 2. IEEE, Taipei (Taiwan), 709–714.
- [10] Luis Omar Colombo-Mendoza, Rafael Valencia-García, Alejandro Rodríguez-González, Giner Alor-Hernández, and José Javier Samper-Zapater. 2015. Recom-Metz: A context-aware knowledge-based mobile recommender system for movie showtimes. *Expert Systems with Applications* 42, 3 (2015), 1202–1222.
- [11] María del Carmen Rodríguez-Hernández and Sergio Ilarri. 2016. Pull-Based Recommendations in Mobile Environments. *Computer Standards & Interfaces* 44 (2016), 185–204.
- [12] María del Carmen Rodríguez-Hernández, Sergio Ilarri, Ramón Hermoso, and Raquel Trillo-Lado. 2017. DataGenCARS: A Generator of Synthetic Data for the Evaluation of Context-Aware Recommendation Systems. *Pervasive and Mobile Computing* 38 (2017), 516–541.
- [13] María del Carmen Rodríguez-Hernández, Sergio Ilarri, Ramón Hermoso, and Raquel Trillo-Lado. 2017. Towards Trajectory-Based Recommendations in Museums: Evaluation of Strategies Using Mixed Synthetic and Real Data. In *Eighth International Conference on Emerging Ubiquitous Systems and Pervasive Networks (EUSPN)*. Elsevier, Lund (Sweden), 1–6.
- [14] Fady Draidí, Esther Pacitti, and Bettina Kemme. 2011. P2Prec: A P2P Recommendation System for Large-Scale Data Sharing. In *Transactions on Large-Scale Data- and Knowledge-Centered Systems III: Special Issue on Data and Knowledge Management in Grid and P2P Systems*, Abdelkader Hameurlain, Josef Küng, and Roland Wagner (Eds.). Springer, Berlin, Heidelberg, 87–116.
- [15] Damianos Gavalas, Vlasios Kasapakis, Charalampos Konstantopoulos, Grammati Pantziou, and Nikolaos Vathis. 2017. Scenic route planning for tourists. *Personal and Ubiquitous Computing* 21, 1 (2017), 137–155.
- [16] Damianos Gavalas, Charalampos Konstantopoulos, Konstantinos Mastakas, and Grammati Pantziou. 2014. Mobile recommender systems in tourism. *Journal of Network and Computer Applications* 39 (2014), 319–333.
- [17] Giuseppe Ghiani, Fabio Paternò, Carmen Santoro, and Lucio Davide Spano. 2009. UbiCicero: A location-aware, multi-device museum guide. *Interacting with Computers* 21, 4 (2009), 288–303.
- [18] Peng Han, Bo Xie, Fan Yang, and Ruimin Shen. 2004. A scalable P2P recommender system based on distributed collaborative filtering. *Expert Systems with Applications* 27, 2 (2004), 203–210.
- [19] Yan Han, Hongzhi Guan, and Jiaying Duan. 2014. Tour Route Multiobjective Optimization Design Based on the Tourist Satisfaction. *Discrete Dynamics in Nature and Society* 2014, Article 603494 (2014), 8 pages.
- [20] Haosheng Huang and Georg Gartner. 2012. Using Context-Aware Collaborative Filtering for POI Recommendations in Mobile Guides. In *Advances in Location-Based Services: Eighth International Symposium on Location-Based Services (Lecture Notes in Geoinformation and Cartography)*. Springer, Vienna (Austria), 131–147.
- [21] Sergio Ilarri, Thierry Delot, and Raquel Trillo-Lado. 2015. A Data Management Perspective on Vehicular Networks. *IEEE Communications Surveys and Tutorials* 17, 4 (Fourthquarter 2015), 2420–2460.
- [22] Sadek Jbara, Tsvi Kuflik, Pnina Soffer, and Olivier Stock. 2007. Context Aware Communication Services in “Active Museums”. In *IEEE International Conference on Software-Science, Technology & Engineering (SwSTE)*, Vol. 2. IEEE, Herzlia (Israel), 127–135.
- [23] Cihan Kaleli and Hüseyin Polat. 2010. P2P collaborative filtering with privacy. *Turkish Journal of Electrical Engineering & Computer Sciences* 18, 1 (2010), 101–116.
- [24] YungMing Li, LienFa Lin, and ChunChih Ho. 2017. A social route recommender mechanism for store shopping support. *Decision Support Systems* 94 (2017), 97–108.
- [25] Sze-Yao Ni, Yu-Chee Tseng, Yuh-Shyan Chen, and Jang-Ping Sheu. 1999. The Broadcast Storm Problem in a Mobile Ad Hoc Network. In *Fifth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*. ACM, New York, NY, USA, 151–162. <https://doi.org/10.1145/313451.313525>
- [26] The Museum of Modern Art (MoMA). 2016. MoMA Collection. (2016). <https://github.com/MuseumofModernArt/collection/tree/v1.12>, [Accessed on August 16, 2017].
- [27] Georgios Pitsilis and Lindsay F. Marshall. 2006. A Trust-enabled P2P Recommender System. In *15th IEEE International Workshops on Enabling Technologies: Infrastructures for Collaborative Enterprises (WETICE)*. IEEE, Manchester (United Kingdom), 59–64.
- [28] Nikolaos Polatidis, Christos K. Georgiadis, Elias Pimenidis, and Emmanouil Stiakakis. 2017. Privacy-preserving recommendations in context-aware mobile environments. *Information and Computer Security* 25, 1 (2017), 62–79.
- [29] Francesco Ricci. 2010. Mobile recommender systems. *Information Technology & Tourism* 12, 3 (2010), 205–231.
- [30] Francesco Ricci, Lior Rokach, Bracha Shapira, and Paul B. Kantor (Eds.). 2011. *Recommender Systems Handbook*. Springer, New York, NY, USA.
- [31] Ankit Rohatgi. 2007. WebPlotDigitizer. (2007). <http://arohatgi.info/WebPlotDigitizer/app>, [Accessed on August 16, 2017].
- [32] Irene Rubino, Jetmir Xhembulla, Andrea Martina, Andrea Bottino, and Giovanni Malnati. 2013. MusA: Using Indoor Positioning and Navigation to Enhance Cultural Experiences in a Museum. *Sensors* 13, 12 (2013), 17445–17471.
- [33] Tuukka Ruotsalo, Krister Haav, Antony Stoyanov, Sylvain Roche, Elena Fani, Romina Delia, Eetu Mäkelä, Tomi Kauppinen, and Eero Hyvönen. 2013. SMART-MUSEUM: A mobile recommender system for the Web of Data. *Web Semantics: Science, Services and Agents on the World Wide Web* 20 (2013), 50–67.
- [34] Taro L. Saito. 2016. SQLite JDBC Driver 3.14.2. (2016). <https://bitbucket.org/xerial/sqlite-jdbc>, [Accessed on August 16, 2017].
- [35] Chih-Yung Tsai, Shuo-Yan. Chou, and Shih-Wei. Lin. 2008. Location-Aware Tour Guide Systems in Museum. In *Collaborative Product and Service Life Cycle Management for a Sustainable World: 15th ISPE International Conference on Concurrent Engineering (CE) (Advanced Concurrent Engineering)*. Springer, London, UK, 349–356.
- [36] Amund Tveit. 2001. Peer-to-peer Based Recommendations for Mobile Commerce. In *First International Workshop on Mobile Commerce (WMC)*. ACM, Rome (Italy), 26–29.
- [37] U. Allemandi & Co. Publishing Ltd. 2016. Visitor Figures 2015 – The Grand Totals: Exhibition and Museum Attendance Numbers Worldwide. *The Art Newspaper* XXV, 278 (April 2016).
- [38] Peng Wang, Jing Yang, and Jian-Pei Zhang. 2016. Protection of Location Privacy Based on Distributed Collaborative Recommendations. *PLOS ONE* 11, 9 (2016), 1–22.
- [39] Wan-Shiou Yang and San-Yih Hwang. 2013. iTravel: A recommender system in mobile peer-to-peer environment. *Journal of Systems and Software* 86, 1 (2013), 12–20.
- [40] Kam Fung Yeung, Yanyan Yang, and David Ndzi. 2010. Context-aware news recommender in mobile hybrid P2P network. In *Second International Conference on Computational Intelligence, Communication Systems and Networks (CICSyN)*. IEEE, Liverpool (United Kingdom), 54–59.
- [41] Yaming Zhang, Haiou Liu, and Shiyong Li. 2011. A Distributed Collaborative Filtering Recommendation Mechanism for Mobile Commerce Based on Cloud Computing. *Journal of Information & Computational Science* 8, 16 (2011), 3883–3891.