Erschienen in: Information Technology & Tourism : ITT ; 10 (2008), 4. - S. 297-311
https://dx.doi.org/10.3727/109830508788403178

Recommendation based on Opportunistic Information Sharing between Tourists

Alexandre de Spindler Moira C. Norrie
Michael Grossniklaus
Institute for Information Systems, ETH Zurich
CH-8092 Zurich, Switzerland

Corresponding Author
Alexandre de Spindler
Haldeneggsteig 4/Weinbergstrasse
CH-8092 Zurich, Switzerland
Phone: +41 44 632 74 16
Fax: +41 44 632 18 92
despindler@inf.ethz.ch

Abstract

We propose a new approach to collaborative filtering in mobile tourist information systems based on spatio-temporal proximity in social contexts. The approach is motivated by a survey of festival visitors confirming that similarities of interests extends beyond events defining specific social contexts. We show how opportunistic information sharing in mobile ad-hoc networks can be used to realise decentralised collaborative filtering appropriate for mobile environments and show its equivalence to existing centralised approaches.

Mobile Information System, Spatio-Temporal Proximity, Copresence, Social Context, Opportunistic Sharing, Ad-hoc Networks, Collaborative Filtering

1 Introduction

Organisers of large arts festivals such as the Edinburgh Fringe are keen to find ways in which users can share ratings and recommendations of events. A typical solution is to provide a section on the festival web site where users can enter reviews. In practice, however, relatively few visitors take the time and effort to write reviews once they have returned home. The other problem is that reading through lots of reviews can be a very time-consuming way of arriving at recommendations taking into account individual user preferences. Other approaches such as the use of sms messages to rate events has the advantage that it is fast and simple to enter the information and can be done immediately after the event. However some experiments with this had a very low take-up which may be due to the costs borne by the user or the lack of awareness of what was a stand-alone service.

Our goal was to develop an approach where users could enter and receive recommendations on the move based on collaborative filtering (CF) techniques used in recommender systems. The ratings and reviews received should be filtered and ranked according to user similarity so that a highly rated event can really be considered as a personalised recommendation. Further, it should be possible to integrate the recommender system into a festival guide available for mobile devices such as the EdFest system (Norrie et al., 2007) to encourage usage. Such mobile systems differ from traditional systems with CF in two ways. First, a mobile guide such as EdFest can deliver context-aware information to tourists without the need to have a network connection and therefore, ideally, we would prefer not to have a central server in the architecture since it significantly increases costs and reduces performance. Second, unlike on-line stores such as Amazon, user and item profiles cannot evolve gradually over time. Visitors to the festival will want a filtering of relevant information and recommendations as soon as they arrive at the festival for what might be a visit of only one or two days.

In this paper, we show how opportunistic information sharing in mobile ad-hoc networks could be used to realise a decentralised approach to collaborative filtering that is equivalent to existing centralised approaches. The underlying assumption is that users who share social contexts have similar interests and this can be used as a basis for filtering recommendations. The idea of using physical copresence as a basis for forming social networks has been investigated in a number of projects. For example, AIDE (Ambient Information Dissemination Environment) (Lawrence, Payne, & Roure, 2006) uses ad-hoc Bluetooth connections between mobile devices to share content such as photos, articles, jokes and upcoming events between users. They propose the use of data mining algorithms to analyse encounters and form social networks, introducing the concept of a *copresence community* as a group of individuals who regularly share the same location at the same time. The TRACE project (Counts & Geraci, 2005) also investigated the use of physical copresence as a basis for forming social networks. To test the assumption that people who attend the same events have similar interests, they recruited users at various social events and then analysed their usage of a web-based system that allowed them to contact other users who attended the same event. The results of this initial field study showed that users generally found the system useful and in some cases it led to users planning future social activities together. Although these projects investigated the use of physical copresence as a means of forming social networks and sharing information, they have not considered how collaborative filtering algorithms could be adapted to base user similarity on shared social contexts. Nor have they developed any general infrastructure to support the opportunistic sharing of information between personal databases in mobile environments. Our goal was to do exactly that and investigate the use of peer-to-peer architectures to allow users to exchange data automatically and unobtrusively based on spatio-temporal proximity.

The issue of information sharing in mobile ad-hoc networks is often seen as the problem of how to ensure that users can access remote data in networks without a fixed topology and with possible disconnections. However, as we have seen above, the ad-hoc nature of establishing network connections between personal mobile devices can be viewed as a means of sharing information *opportunistically* among members of a user community based on spatio-temporal proximity. In the context of a festival, this would mean that users would receive ratings and reviews from users attending the same event which should be more useful than those from arbitrary users. The higher the frequency of encounters between two users, the greater the similarity between these users is likely to be and hence their respective recommendations should be more highly regarded.

We begin in section 2.1 with a discussion of requirements for collaborative filtering in mobile settings. Section 2.2 reports on a field study carried out at the Edinburgh Festivals in 2006 to test the assumption that physical copresence in social contexts can be used as a measure of user similarity in collaborative filtering. We then present the details of our collaborative filtering algorithm based on opportunistic information sharing in peer-to-peer environments in section 2.3. In section 2.4, we prove the correctness of our approach by showing that it is equivalent to existing collaborative filtering techniques. Finally, we discuss some issues of our approach in section 2.5 and give concluding remarks in section 3.

2 Collaborative Filtering in Mobile Settings

2.1 Requirements

The motivation for our approach to collaborative filtering in mobile settings came from our experiences of developing the EdFest system, which was a prototype mobile information system for visitors to the Edinburgh festivals based around a set of interactive paper documents (Norrie et al., 2007). Users could interact with the documents using a digital pen and output was delivered through an audio channel. An example of a brochure entry is shown in figure 1. By touching any of the pictograms with the digital pen, a user will activate the associated information service.

A key feature of our system is to provide an easy means for users to input and access ratings and reviews. A rating can be entered by touching the appropriate pictogram for the range 1–5 in the top right corner of an event entry. The average rating can be accessed by by touching the pictogram to the left of 'Rating'. Reviews could be entered as handwritten comments in the back of the brochure and accessed by touching the pictogram with the speech bubbles under the event title.

Like many other mobile festival guides, most of the information in the

Figure 1: Interactive festival brochure entry

EdFest guide is relatively static and could be downloaded and stored in a mobile computing device before the festival visit. However, the ratings and reviews are dynamic and require access to a central server. Further, simply providing access to an average of all ratings and all reviews does not take into account the fact that users may have quite different tastes. Our goal was to find a way in which we could improve recommendations by taking user similarities into account and, at the same time, improve performance by removing the need for a central server.

Collaborative filtering systems have been developed to help users make a choice among unknown alternatives by trying to find other users with similar interests and tastes. The incorporation of other people's opinions in a decision process requires some form of social network where opinions can be selectively retrieved and combined. In the scope of this work, we simplify the main task of CF to inferring an opinion for a requesting user about a target item unknown to them. We define an opinion as a rating value which is low to express a bad opinion and high otherwise. In order to generate recommendations and thus fulfil the claim of CF, we can think of ranking all items according to their ratings and selecting the highly ranked ones.

Recommender systems based on collaborative filtering (CF) have become well-known through their use in on-line stores. The underlying assumption is that users who bought the same items in the past are likely to do so in the future. The most fundamental query to a CF system consists of inferring an opinion of a user about an item (Aggarwal, Wolf, Wu, & Yu, 1999). Most CF approaches share the main steps in inferring an opinion. These steps are the assessment of similarities between objects, i.e. users or items, the selection of similar objects and the aggregation of opinions about a set of objects. We will briefly describe the variety of CF techniques that have been proposed before discussing the requirements of CF techniques suited to mobile environments.

One of the first approaches developed was user-based CF (Resnick, Iacovou, Suchak, Bergstrom, & Riedl, 1994), in which the fundamental query is processed by selecting a set of users similar to the given one and aggregating their opinions about the specified item. The similarity between users is measured in terms of the extent to which their opinions about items correlate. User-based CF has been deployed in a wide variety of application domains such as usenet news (Konstan et al., 1997), music (Shardanand & Maes, 1995), video (Hill, Stead, Rosenstein, & Furnas, 1995) and web page (Terveen, Hill, Amento, McDonald, & Creter, 1997) recommendations as well as e-commerce (Schafer, Konstan, & Riedi, 1999). As a result of all the experience gained, three main shortcomings of user-based CF are frequently given, namely the issues of sparsity, scalability and cold starts. Firstly, the number of items a user has provided an opinion on is typically sparse compared to the total number of items. If the number of commonly rated items is small, the similarity will be inaccurate and so will the prediction. Secondly, the complexity of selecting a set of similar users grows with the number of users and items as $O(|users| \times |items|)$ leading to problems of scalability. Thirdly, new users will have expressed opinions about only a few items leading to the same problems of sparsity. Similarly, when a new item is introduced and it has not been rated by any users, no opinion about it can be inferred.

A number of approaches address the shortcomings of user-based CF while retaining its advantages. Sarwar, Karypis, Konstan, & Reidl (2001) introduced the idea of item-based CF where the fundamental query is processed by selecting a set of items similar to the one for which an opinion is to be inferred and aggregating all opinions about them. User- and item-based CF are the most well known representatives of so called memory-based approaches. Memory-based refers to the technique of performing filtering based on the raw data, a set of tuples each containing a rating user and an item rated with a numerical value expressing the opinion of the user about the item. This differs to model-based approaches that are characterised by computing an intermediate representation of the set of the tuples such as clusters (Ungar & Foster, 1998), probability distribution functions (Breese, Heckerman, & Kadie, 1998) or singular value decompositions (Zhang, Wang, Ford, Makedon, & Pearlman, 2005). Model-based approaches effectively resolve the sparsity issue and render predictions more efficient and supposedly accurate.

Hybrid filtering systems combine elements from memory- and model-based CF (Pennock, Horvitz, Lawrence, & Giles, 2000) as well as content-based filtering (Polcicova, Slovak, & Navrat, 2000). While content-based filtering comes with the burden of extracting intrinsic properties of unstructured information or natural language content, it does not suffer from sparse data. It has therefore proven successful for resolving the sparsity and cold start issues.

Most collaborative filtering systems have been designed to be deployed in client-server architectures whereas only a few approaches (Wang, Pouwelse, Lagendijk, & Reinders, 2006; Miller, Konstan, & Riedl, 2004; Tveit, 2001) have tackled the challenges of decentralised environments. Distributed filtering research has mainly been concerned with the availability of data on client devices where network connectivity cannot be guaranteed and opinions need to be predicted. In addition to challenges of distribution, mobile environments impose particular requirements to collaborative protocols. Devices must be portable which restricts their size, weight and electrical power capacity, and these restrictions imply limited computational power and human computer interface facilities. Despite the advantages of model-based CF in comparison with memory-based approaches, computing the intermediary representation emerges as a new bottleneck, in particular with regard to the limited computational power available on mobile devices. In the case of content-based approaches, the limited interaction facilities offered by mobile devices make it impractical to have users specifying properties of content and automatic extraction requires intensive computations. Further, although wireless connectivity is increasingly available within restricted areas such as restaurants and airports as well as public areas by means of 3G networks, area-wide connectivity is still bound to expensive communication costs, high power consumption and prone to disconnections. In contrast, devices may connect to each other in an ad-hoc peer-to-peer fashion based on short range connectivity technology such as Wi-Fi and Bluetooth.

Consequently, a CF protocol for mobile environments must respect the following requirements. All computation and storage must be decentralised since a connection to a central server may not be available. Due to restricted computational and storage capacities of mobile devices, local computation must be kept simple and the required data small. Ideally, the protocol should rely on ad-hoc peer-to-peer connections only. This transient connectivity requires data exchange to be short and to consume little bandwidth, in particular for the case of Bluetooth technology. Additionally, the protocol must be delay tolerant since other peers may not always be available. Finally, since mobile devices typically feature reduced interaction facilities, user interaction should be minimal.

As we will show in the following sections, the results of a study indicate a correlation between physical proximity and user similarity which can be used to reduce the computing costs of CF as well as rendering CF suitable for ad-hoc connectivity available in mobile environments. Users who are close enough to each other to get connected using short range network technologies and stay in proximity for a sufficient amount of time tend to share similarities that can be used instead of, or in addition to, the similarities computed based on rating tuples, intermediary representations or profiles.

2.2 Festival Field Study

If users tend to be in the same social environment, at the same time, then they typically share some social preferences. For example, if two users attend the same music concert, it is likely that they have similar musical tastes. Initial studies such as those carried out in the TRACE project (Counts & Geraci, 2005), support the assumption that the notion of *shared social contexts* can be exploited to establish a similarity relationship between users. In order to show the relationship between spatio-temporal proximity and user-similarity, we conducted a survey where we assessed the tastes and interests of participants of a large-scale festival in Edinburgh. The Edinburgh festival is in fact a collection of festivals including the Fringe Festival, a Book Festival, a Film Festival and the Military Tattoo. The Fringe is the world's largest arts festival consisting of various categories of events including comedy, music, theatre and dance. It runs over a four week period with over 300 venues, 1,800 events and 28,000 performances. Of special interest was to see whether user similarities could be extended beyond the specific category of event defined by a social context. For example, if two users attend the same play, is it also likely that they have similar preferences for music, books, films etc?

We interviewed visitors at seven venues associated with different categories of events. A minimum of 30 visitors were interviewed at each of the venues and we also interviewed people at the main railway station as an example of a general public place. The questionnaire used in all of these interviews was split into two parts. While the first part allowed festival visitors to express their opinion on a set of books, films and music albums, the second part was used to assess which festivals and event categories they had visited. The sets of books, films and albums in the first part were methodically selected based on a combination of awards, public consumption and genre for the year 2003. For each of the book, film and album items, participants could choose one out of six answers. Three of these answers were applicable in case the participant had consumed the item, i.e. had read the book, seen the film or listened to the album. One of the other three answers could be chosen if the item was unknown. In both cases, three answers were available to express a positive, neutral or negative opinion. In the second part of the questionnaire, participants were asked to answer yes-or-no questions to state whether they had visited, were currently visiting or intended to visit one of the various festivals as well as different categories of events such as theatre, comedy and visual art.

A standard chi square (χ^2) test was used to investigate the relationship between the time and the location at which participants were interviewed as well as the frequency of each possible answer for a particular question. While a chi square test allows a relationship between two variables to be measured, no conclusions can be made about the direction of the relationship. However, since we are only interested in testing whether people in spatio-temporal proximity tend to give similar answers to our questionnaire, it is of no relevance to us whether the location causes people to answer similarly or if people's similarity causes them to be in proximity.

First, we identified a relationship among people in proximity and which events they were attending, apart from the one where they had been interviewed. To do so, we used the data collected in the second part of the questionnaire where participants were asked about all festivals and the event categories dance, musical, opera, classical music concert, theatre, comedy, talks and visual art. The answers to all of them except for dance and musical events differed significantly. For example, participants interviewed at the Military Tattoo, Film Festival, Book Festival and Fringe Festival showed significantly different frequency distributions of yes-or-no answers when asked about whether they had visited, were currently visiting or intended to visit comedy events. Figure 2 plots the distributions per location which differ on a significance level of one percent.

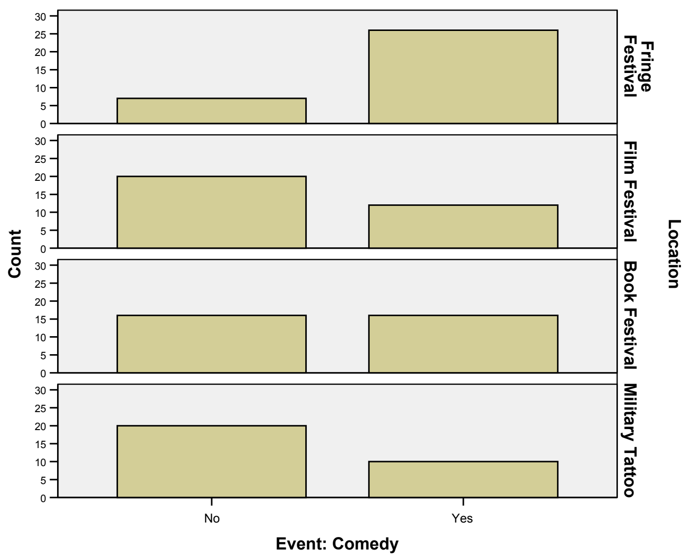

Figure 2: Distributions of answers about comedy events per location

These results led us to conclude that users consuming a particular item also share similar preferences about other items apart from the one they are currently consuming. An intuitive interpretation of this result with respect to collaborative filtering would be that a user collecting recommendations while consuming an item should receive information about other categories of items that they tend to consume themselves. Therefore these recommendations would be valuable for them if they had not yet consumed these items.

The answers to the book, album and film items collected in the first part of the questionnaire were analysed in two steps. First, we aggregated all answers such that we only considered whether the participants had consumed an item or not. The results showed that similarity of tastes can be generalised to other types of items such as books, films and music. However, certain individual items showed a clearer relationship than others and the results were the least clear in the case of music albums. It proved to be much more difficult to find a suitable set of specific albums to act as classifiers, partly because people were less familiar with specific albums than the associated artists. However, people in proximity tended to know about the same items and if they would exchange ratings they would give each other appropriate recommendations. In figure 3, we plot the distributions of answers to an example film item per location which differ on a significance level of one percent.

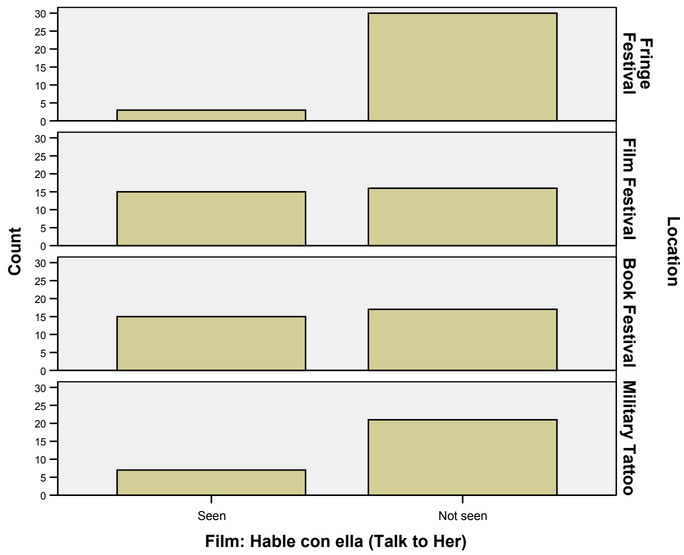

Figure 3: Distributions of answers about a film per location

Finally, we also analysed the opinions of the participants. For this purpose, the answers for each book, album and film item were aggregated such that we only considered the positive, neutral or negative opinion regardless of whether they knew about the item or not. Although not all items showed a significant relationship, the results indicate that people in proximity tend to share similar opinions about items not related to the one they are currently consuming. Figure 4 shows the distributions of opinions about an example book item which differ on a significance level of one percent.

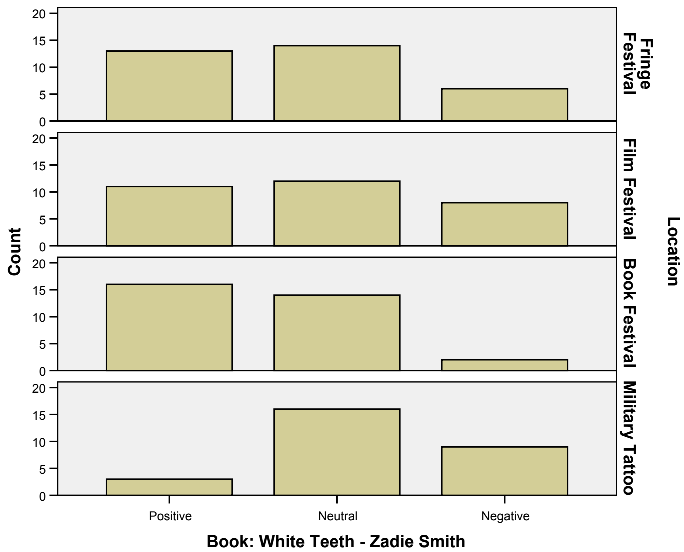

Figure 4: Distributions of answers about a book per location

While some items that were chosen when the questionnaire was designed showed no statistical significance, the results of this survey, on the whole, support our hypothesis that people sharing a social context simultaneously tend to have similar tastes and interests.

2.3 Spatio-Temporal Collaborative Filtering

The application domain of CF contains users consuming items and expressing opinions about these items. Based on these, a collaborative filtering system predicts their opinion about items unknown to them. Opinions are tuples of the form $(user, item, value)$ containing a user expressing the opinion, the item subject to judgement and a value representing the opinion. This is a general form of the application domain, defining the concepts and the form of data expected by most of existing CF approaches. In this work we were keen on sticking to general forms as it would allow for other approaches to be easily integrated with the one presented here which further increases the quality of predictions. Moreover, the implementation of our approach has been parametrised in order to being able to experiment with critical aspects.

Opinion tuples can be seen as a directed weighted edge in a graph, pointing from a user node to an item node and weighted with a rating value. Thus, a set of tuples defines a directed graph $G = (U \cup I, E)$ where U is the set of nodes representing users, I the set of item nodes and E the set of directed weighted edges pointing from nodes in U to nodes in I.

User-based CF processes a fundamental query by first computing similarities among users and selecting those judged to be similar. Then the ratings of target items by these users are aggregated. In order to include user similarities, we augment the previously defined graph with undirected edges connecting two users and weighted with their similarities. Thus, the set of edges E is now composed of $E_r \cup E_s$ where E_r contains the rating edges and E_s the similarity edges. As proposed by Mirza et al. (Mirza, Keller, & Ramakrishnan, 2003), $G_s = (U, E_s)$ represents a *social network graph* while $G_r = (U \cup I, E_r)$ refers to the *rating graph*.

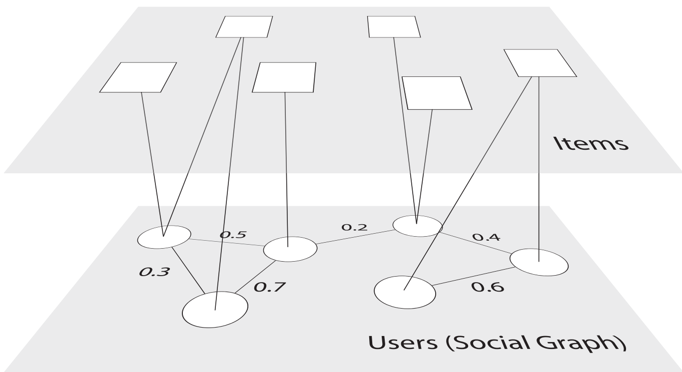

Figure 5: Social- and rating graph

Figure 5 shows an example graph composed of a social and rating graph. The vertices on the bottom layer represent users and the ones on the top layer items. Edges connecting users are weighted with the similarity of the adjacent users. For clarity, we omit the weights of the edges connecting a user to an item representing the rating value.

Normally the first two stages of user-based CF are concerned with the computation of similarities between users and the selection of those most similar. However, in our approach, this selection of users is performed implicitly and in the absence of any prior similarity computations. We introduce the concept of *spatio-temporal proximity* which forms the basis for our selection of similar users.

In the case of social contexts formed around consumable items, user consumption of an item means that their location matches the location of the item for a specific period of time. Some items such as restaurants or bars can be consumed at any time within predefined opening hours and the duration of consumption can be anything from the time to drink a glass of wine up to eating a dinner. In contrast, items such as comedy shows or theatre plays can be consumed only during a specific time period and the duration is usually well defined. We will refer to these two kinds of items as *location* and *event* items, respectively. Note that event items may happen only once or be repeated periodically.

All items have in common the fact that if users meet while consuming them, they stay in each other's vicinity for longer than if they would pass each other in the street by chance. As has been pointed out by Lawrence et al. (2006) and supported by our field study, users who consume the same items share similar properties in terms of interest and taste. Our selection of similar users takes advantage of the implication that users who find themselves at the same location because they are consuming the same item tend to be similar.

This similarity can be derived from spatio-temporal proximity which is equivalent to the notion of copresence. If multiple users are consuming a particular item, their locations will match the location of the item. Thus, their spatial proximity will not exceed an item-specific boundary. Also, if users are consuming an item simultaneously, the period of time during which they are consuming it will overlap. This overlap is a result of temporal proximity. We conclude that if users are in spatio-temporal proximity, they are consuming the same item simultaneously and thus share similarity properties. Most importantly, the similarity implies that they have been in each other's vicinity at some point in time.

The history of item consumption of a particular user u_a can be regarded as a set of *item consumption tuples* of the form $(loc_i, [t_k, t_l])$ where each tuple contains two entries. The first entry identifies a location loc_i particular to an item. This location represents an area in which the item can be consumed. The second one delimits a period of time $[t_k, t_l]$ during which the item was consumed. Consequently, the history $H(u_a)$ of a user u_a can be written as

$$H(u_a) = \{(loc_1, [t_1, t_2]), (loc_2, [t_3, t_4]), \ldots\}$$

The condition for item consumption tuples to be equal is

$$(loc_i, [t_k, t_l]) = (loc_j, [t_m, t_n]) \iff (loc_i = loc_j) \wedge ([t_k, t_l] \cap_t [t_m, t_n] \geq p)$$

where we define $\cap_t$ as a temporal intersection of two time periods. The condition $[t_k, t_l] \cap_t [t_m, t_n] \geq p$ holds if the time periods overlap for a duration of at least p. The first component $(loc_i = loc_j)$ accounts for spatial proximity while the temporal intersection accounts for temporal proximity.

The user similarity $P_{loc,t}$ resulting from spatio-temporal proximity between two users u_a and u_b can be expressed as

$$P^{\mathbb{N}}_{loc,t}(u_a, u_b) = \begin{cases} 1 & \text{if } H(u_a) \cap H(u_b) \neq \emptyset \\ 0 & \text{else} \end{cases}$$

This is a binary similarity measure in the sense that users are evaluated to be similar only if they have at least one tuple of their consumption history in common. We use $P^{\mathbb{N}}_{loc,t}(u_a, u_b)$ as a condition for the users $u_a \in U$ and $u_b \in U$ to be connected by a similarity edge $(u_a, u_b) \in E_s$ in the social graph G_s. The resulting social graph corresponds to a copresence community used by Lawrence et al. (2006) to disseminate information since spatio-temporal proximity is a necessary and sufficient condition for users to have their devices connected.

We can refine this similarity taking into consideration the level of spatio-temporal proximity among users. Based on the fact that users consuming the same items are similar, it is obvious that the more often users consume the same item, the more similar they are. This calls for a continuous similarity measure $P^{\mathbb{R}}_{loc,t}$ that takes into account the number of common simultaneous item consumptions as opposed to the binary measure proposed before.

$$P^{\mathbb{R}}_{loc,t}(u_a, u_b) = \begin{cases} \frac{|H(u_a) \cap H(u_b)|}{max(|H(u_a)|, |H(u_b)|)} & \text{if } H(u_a) \neq \emptyset \\ 0 & \text{else} \end{cases}$$

This measure allows us to assign a weight to a similarity edge created based on the binary measure. Note that if it evaluates to zero, the respective users are not connected in the graph, while it never evaluates to zero if they are connected.

We now describe our CF approach in terms of a formal description of the algorithm running on a single mobile device as shown in figure 6. For this discussion, we assume the existence of three library functions. WAIT(p) causes the algorithm to pause for a time period of p, TRANSMIT($Peer, M$) transmits a set of edges M to a remote peer $Peer$. This transmission will be translated to a call of the function RECEIVE(M) on the remote peer where M corresponds to the second argument of the transmission function. INCREASE-WEIGHT($Peer$) retrieves the edge $(u_{local}, u_{remote}) \in E_s$ where u_{remote} denotes the user node representing the argument $Peer$ and increases its weight in order to update the respective continuous proximity value.

```
MAIN-LOOP()
N ← ∅
while run = ⊤
do N_current ← SCAN()
   N_new ← N_current − N
   for ∀ Peer ∈ N_new
   do SEND(Peer)
      INCREASE-WEIGHT(Peer)
   N ← N_current
```

```
SEND(Peer)
M ← ∅
for ∀ (u_local, i) ∈ E_r
do M ← M ∪ {(u_local, i)}
WAIT(p)
TRANSMIT(Peer, M)
```

```
RECEIVE(M)
for ∀ (u_remote, i) ∈ M
do E_r ← E_r ∪ {(u_remote, i)}
```

Figure 6: Collaborative filtering algorithm

While a peer is active, i.e. $run = \top$, the main loop simply scans the environment periodically and maintains a set N of peers in the vicinity. For every remote peer $Peer$ in the vicinity, the method SEND($Peer$) is called to send all ratings made by the local user to the remote peer. This method runs as a thread per remote peer in order to be non-blocking. Note that these ratings will only be sent after a delay of length p, the parameter introduced above to determine the equality of two rating consumption tuples. If the remote peer has left the vicinity of the local peer during this time period, the tuples will not be sent by the $\textsc{Transmit}(Peer, M)$ function to avoid exchanges during a transient encounter. Once the rating tuples have been sent to all new peers in the vicinity, the set of peers in the vicinity is updated to remove peers that have left. Whenever a local peer receives a set of tuples from a remote peer, $\textsc{Receive}(M)$ is called and these tuples are added to the set of tuples stored locally.

Finally, rating values from similar users about the target item are aggregated. The most common approach is to compute the average. To do so, we select all incoming edges of the node representing the target item and compute the average of their weights. We also take into account the degree of similarity as expressed by the continuous proximity measure. $P^{\mathbb{R}}_{loc,t}(u_a, u_b)$ establishes a ranking of the users according to their similarity to the user denoted by the first argument. A user u_a is more similar to a user u_b than to another user u_c if $P^{\mathbb{R}}_{loc,t}(u_a, u_b) > P^{\mathbb{R}}_{loc,t}(u_a, u_c)$. Consequently, if we are to predict a rating value for a requesting user u_r about a target item it_t, we compute the average of the rating values contained in G_r, each weighted with the respective edge weights in G_s. When computing this weighted average, we only need the continuous proximity values for the rating user to all other users in the local graph. The similarity between other users does not affect the aggregation and thus no continuous proximity information needs to be passed on when ratings are exchanged.

2.4 Equivalence to Existing Algorithms

As explained in the previous section, users of our recommender system exchange tuples when they are in spatio-temporal proximity. Each user maintains a graph G_{local} where the nodes in U represent users previously met and the nodes in I represent all items rated by these users or the local user. In this section, we first explain why such a local graph is sufficient to perform user-based collaborative filtering. Secondly, we show that the resulting algorithm resolves scalability issues for which user-based approaches have frequently been criticised.

We first look at a simple form of traditional user-based CF where rating values are set to 1 if a user has consumed an item and 0 otherwise. For example, the Amazon online store interprets the purchase of an item as an expression of a binary opinion about it. Thus, each user is represented by a binary vector containing entries for all items. A server maintains the set of user vectors based on which ratings are predicted. The similarity between two users is computed as the number of vector entries both have set to 1. The prediction is the result of aggregating the ratings of all users about the target item, each weighted with the similarity between the requesting and rating user. Consequently, the prediction is based on the set of users that have consumed at least one item which the requesting user has also consumed. All other users are not included because their ratings are weighted with a zero-valued similarity.

A user vector is a set of rating tuples where the user entry contains the represented user. The tuples stored on the server define a graph G_{global} which, in contrast to a local graph, includes all participating users and items consumed by any user. Therefore, a local graph is a subgraph of the global graph while the global graph is a union of all local graphs. In fact, a local graph belonging to a particular user u_i can be extracted from the global graph as follows. We use superscript notations g and l to indicate that a node or edge set belongs to the global or local graph, respectively. An edge is denoted as (p, q) where p and q are the adjacent nodes. Finally, $w_{(p,q)}$ refers to the weight of an edge (p, q).

$$U^l = \{u \mid u \in U^g \wedge (u_i, u) \in E_s^g \wedge w_{(u_i,u)} > 0\} \cup u_i \quad (1)$$

$$I^l = \{i \mid i \in I^g \wedge (u, i) \in E_r^g \wedge u \in U^l\} \quad (2)$$

$$E_r^l = \{(u, i) \mid (u, i) \in E_r^g \wedge u \in U^l \wedge i \in I^l\} \quad (3)$$

$$E_s^l = \{(u_i, u) \mid (u_i, u) \in E_s^g\} \quad (4)$$

Equation 1 states that we take all users in U^g which are connected to u_i by a similarity edge with a weight greater than zero. Equation 2 selects all items that are connected to a user selected in equation 1. Equation 3 selects all rating edges whose adjacent user and item have been selected by the previous two equations. Finally, equation 4 accounts for the fact that similarity edges are not exchanged. Thus, only similarity edges between u_i and the other users are selected. In figure 7, we highlight the local graph as part of the global graph. Nodes and edges not belonging to the local graph are drawn with a dashed line.

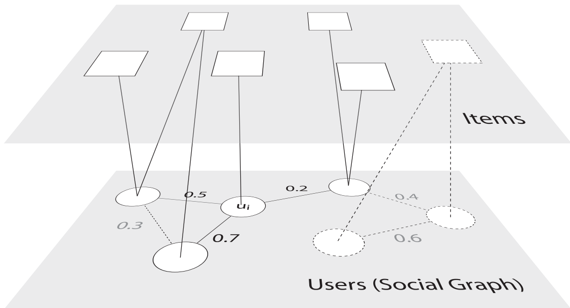

Figure 7: Local graph as a subgraph of the global graph

Simple traditional CF outlined above predicts a rating for a requesting user u_r about a target item i_t as

$$\frac{1}{|(u, i_t) \in E_r^g : (u_r, u) \in E_s^g|} \sum_{(u,i_t) \in E_r^g} w_{(u,i_t)} \cdot w_{(u_r,u)} \tag{5}$$

where the aggregation is a weighted average of the ratings. Now we want to show that all rating and social edges included in the aggregation also exist in the local graph belonging to u_r. The underlying intuition is that users from whom ratings are aggregated have in common the fact that they consumed items also consumed by the requesting user. Hence, if users exchange their own ratings whenever they consume the same item, the set of users from whom opinions are collected and thus are available in the local graph is equivalent to the set of users selected in the global graph by traditional user-based CF. In order to prove this equivalence, we have to show that all rating and social edges included in the sum in equation 5 also exist in the local graph. This is obvious for the social edges because all edges $w_{(u_r,u)} \in E_s^g$ have been selected by equation 4 and, since we are considering the local graph belonging to u_r, it holds that $u_i = u_r$. In order to simplify this proof of equivalence, we can now leave out the weighting of each rating. Therefore we rewrite equation 5 as

$$\frac{1}{|(u, i_t) \in E_r^g : (u_r, u) \in E_s^g|} \sum_{(u,i_t) \in E_r : (u_r,u) \in E_s^g \wedge w_{(u_r,u)} > 0} w_{(u,i_t)} \tag{6}$$

where the condition of the sum ensures rating edges are included only from rating users that have a non-zero similarity to the requesting user. Now it is apparent that the rating edges included in the aggregation are also contained in the local graph since E_r^l is extracted from E_r^g by applying equation 1, 2 and 3 consecutively, while the selection criteria of the sum is equivalent to equation 1. Figure 8 highlights the nodes and edges used for the rating inference. By comparing it with figure 7, it becomes apparent that the same nodes and edges are available in a local graph.

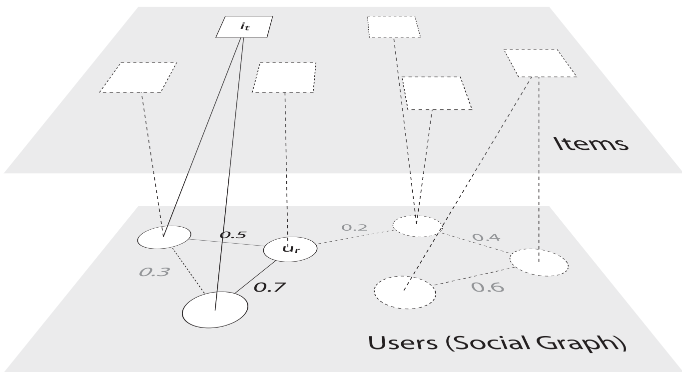

Figure 8: Nodes and edges in the global graph used for rating inference

Since a local graph contains all edges used by traditional user-based CF for rating predictions based on the global graph, we might as well use the local graph and still obtain the same results. However, when using the local graph, there is no need for a central server to store all user vectors, either to compute similarities or to select similar users. The local graph owned by a particular user can be seen as a selection of the users, items, rating- and social edges from the global graph which are relevant to all predictions to be made for the owning user. Hence, making a prediction simply consists of aggregating all rating edges pointing to the target item, each weighted with the respective social edge weight. User-based CF has often been criticised for the fact that it does not scale well. However, it is the similarity computation and selection of similar users that forms the bottleneck in traditional CF. Since our approach bypasses this bottleneck, we have successfully eliminated these performance issues while the resulting predictions are equivalent.

2.5 Discussion and Issues

In order to experiment with our CF approach, we are currently developing a simulation framework allowing the dissemination of opinions to be simulated and observed. In contrast to real world experiments where few users would be equipped with devices implementing our approach, a simulation allows our models to be tested on a larger scale. The simulation framework features locations where items can be consumed and where users choosing to consume particular items exchange opinions. The simulation can be set up with single or multiple data sets such as film and book ratings (Collaborative Filtering Resources, 2007) which determine the items offered at locations and chosen by users in a probabilistic manner. Simulated users are equipped with a virtual system implementing the CF algorithm we have presented. The simulation will therefore allow us to investigate the local graphs defined by the opinions collected.

In order for collaborative filtering based on opportunistic information sharing to yield accurate predictions, it is important that users only exchange ratings when they are consuming the same items. In section 2.3, we explained how we can control the exchange by delaying the transmission of rating tuples. This delay corresponds to the parameter p defined in that section which determines the duration of temporal proximity required for consumption tuples to be equal. If the value of p is too small, users are assumed to have consumed the same item simultaneously as a result of a transient encounter. In contrast, if it is too large, users who actually do consume the same item simultaneously are not recognised as doing so. Our simulation framework allows us to experiment with this parameter in order to find values that rule out transient encounters while accounting for overlapping periods of consumption. We also propose to maintain a collection of locations defining where users exchange ratings. Thus, connectivity can be inhibited when users are in proximity outside the locations specified in this collection. Furthermore, the value of p can be a function of the location in order to account for activity-dependent periods of time signifying simultaneous consumption.

It is also important that users do not exchange ratings multiple times while consuming a single item. This would illegitimately increase their assumed similarity which would result in inaccurate predictions. In order to avoid this, the main loop of our algorithm presented in section 2.3 keeps track of encounters while the users do not significantly change their location. As the realisation of such a location tracker is straightforward, we do not present the details of its implementation. Also, this tracker can take into account the collection of predefined locations introduced above.

We have presented a filtering technique that results from users exchanging ratings when they are in spatio-temporal proximity. Although it can be used on its own, existing CF techniques can still be used to further increase the accuracy of predictions or to grow a set of increasingly similar users based on which predictions are made. Due to the pre-selection based on spatio-temporal proximity, the local graph is significantly smaller than the global graph used in traditional, server-based CF. Thus, our approach effectively addresses the scalability aspect of existing CF approaches while allowing them to be adopted without adaptation efforts.

The graph containing all rating tuples recorded by the local user and received from remote users forms a set of tuples available for rating prediction. This set can also be seen as a pre-selection of similar users which can serve as a basis for further refinement using existing memory- or model-based CF algorithms. Spatio-temporal proximity still contributes to the efficiency of such algorithms since it reduces the set of rating tuples on which computations are performed.

Consequently, the implementation of our CF approach within the simulation framework allows the arbitrary combination of CF algorithms based on a composite pattern, therefore allowing experimentation with various approaches and ways of aggregating each prediction. This serves to further increase the accuracy of predictions. Additionally, similarity computations performed locally, on top of the filtering carried out by our algorithm, prevent the set of tuples from growing infinitely as users move in space and time. Rating tuples from least similar users can be removed periodically which naturally leads to more accurate predictions.

In section 2.3, we distinguish between location and event items. While the consumption of non-periodic event items implies spatio-temporal proximity, users consuming location or periodic event items do not necessarily consume them simultaneously. The fact that users do consume them simultaneously can be regarded as an additional similarity feature thus further increasing the accuracy of predictions. However, users consuming the same location item or attending the same periodic event at different times still share a similarity in terms of interest and taste, even though there is no physical copresence. We therefore want to investigate how we can extend our CF approach to allow for ratings and reviews exchanged between users in spatial, but not temporal, proximity. The same devices as the one carried by the users could be installed at locations where location or periodic event items are consumed. These location-bound devices would act as stationary users, thereby allowing users to exchange opinions despite the lack of temporal proximity. There would be no further issues to consider in the case of location items. However, if we are dealing with periodic event items, we have to allow for the fact that different events may take place at the same location. For example, a movie theatre may show a particular film early in the evening and another one later in the night, both on a daily basis. Obviously, users consuming one of the films should not exchange opinions with users consuming the other one. In order to account for this, the stationary user device would have to be extended to distinguish between opinions exchanged during particular events. Since this device is attached to a location, it would not be difficult to provide it with the information about what event item is currently being consumed.

3 Conclusions

We have presented a technique for user-based collaborative filtering that exploits an opportunistic mode of information sharing resulting from ad-hoc peer-to-peer networking. Only users in spatio-temporal proximity are able to exchange ratings and we have shown how this provides a natural filtering based on social contexts. The resulting selection of similar users renders the computation of similarities and selection of most similar users unnecessary which resolves sparsity and scalability issues frequently associated with user-based collaborative filtering.

We have investigated the similarity measure on which our approach is based by statistically evaluating the results of a survey conducted on visitors of an international arts festival. The results support our hypothesis that people sharing a location simultaneously tend to have similar tastes and interests. We have presented the algorithm performed by mobile devices that users carry around while consuming items. We have shown that rating inference, the fundamental query to a CF system, as performed in our approach is based on the same data as existing centralised CF approaches and therefore the filtering effects are equivalent.

Acknowledgements

We would like to thank Andrei Vancea, Milivoje Petrovic and Mahir Yildirim for their help in conducting the user survey in Edinburgh.

Biographical Note

Alexandre de Spindler is a research assistant in the Global Information Systems research group at ETH Zurich. He received a Diploma (M.Sc.) in Computer Science from ETH Zurich in 2005 and is currently working on his Ph.D. His research interests include architectures, frameworks and platforms for collaborative and mobile information environments as well as mobile and personal information systems.

Moira C. Norrie is a Professor at ETH Zurich where she is head of the Institute for Information Systems and leads the Global Information Systems research group. Her research interests include object-oriented models and systems for data management, web engineering, mobile and personal information systems and interactive paper as a medium for integrating printed and digital information.

Dr. Michael Grossniklaus studied computer science at ETH Zurich where he obtained a diploma in 2001 and a PhD in 2007. Currently, he is working as a senior researcher in the Institute for Information Systems at ETH Zurich. His main research interest is empowering information systems for context-aware data management and delivery in the domain of ubiquitous and mobile computing as well as web engineering.

References

Aggarwal, C. C., Wolf, J. L., Wu, K.-L., & Yu, P. S. (1999). Horting Hatches an Egg: A new Graph-Theoretic Approach to Collaborative Filtering. In *Proceedings of ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Diego, California, United States* (pp. 201–212).

Breese, J. S., Heckerman, D., & Kadie, C. (1998). Empirical Analysis of Predictive Algorithms for Collaborative Filtering. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence* (pp. 43–52).

Collaborative Filtering Resources. (2007). http://ict.ewi.tudelft.nl/˜jun/CollaborativeFiltering.html#dataset.

Counts, S., & Geraci, J. (2005). Incorporating Physical Co-presence at Events into Digital Social Networking. In *Extended Abstracts on Human Factors in Computing Systems, Portland, OR, USA* (pp. 1308–1311).

Hill, W., Stead, L., Rosenstein, M., & Furnas, G. (1995). Recommending and Evaluating Choices in a Virtual Community of Use. In *Proceedings of International Conference on Human Factors in Computing Systems, Denver, Colorado, United States* (pp. 194–201).

Konstan, J. A., Miller, B. N., Maltz, D., Herlocker, J. L., Gordon, L. R., & Riedl, J. (1997). GroupLens: Applying Collaborative Filtering to Usenet News. *Commun. ACM*, *40*(3), 77–87.

Lawrence, J., Payne, T. R., & Roure, D. D. (2006, june). Co-presence Communities: Using Pervasive Computing to Support Weak Social Networks. In *Proceedings of International Workshop on Distributed and Mobile Collaboration, Manchester, UK.*

Miller, B. N., Konstan, J. A., & Riedl, J. (2004). PocketLens: Toward a Personal Recommender System. *ACM Trans. Inf. Syst.*, *22*(3), 437–476.

Mirza, B. J., Keller, B. J., & Ramakrishnan, N. (2003). Studying Recommendation Algorithms by Graph Analysis. *J. Intell. Inf. Syst.*, *20*(2), 131–160.

Norrie, M. C., Signer, B., Grossniklaus, M., Belotti, R., Decurtins, C., &

Weibel, N. (2007). Context-aware platform for mobile data management. *Wireless Networks*, *13*(6), 855-870.

Pennock, D. M., Horvitz, E., Lawrence, S., & Giles, C. L. (2000). Collaborative Filtering by Personality Diagnosis: A Hybrid Memory and Model-Based Approach. In *Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence* (pp. 473–480).

Polcicova, G., Slovak, R., & Navrat, P. (2000). Combining Content-Based and Collaborative Filtering. In *Symposium on Advances in Databses ind Information Systems* (p. 118-127).

Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., & Riedl, J. (1994). GroupLens: An Open Architecture for Collaborative Filtering of Netnews. In *Proceedings of ACM Conference on Computer Supported Cooperative Work, Chapel Hill, North Carolina, United States* (pp. 175–186).

Sarwar, B., Karypis, G., Konstan, J., & Reidl, J. (2001). Item-based Collaborative Filtering Recommendation Algorithms. In *Proceedings of International Conference on World Wide Web, Hong Kong, Hong Kong* (pp. 285–295).

Schafer, J. B., Konstan, J., & Riedi, J. (1999). Recommender Systems in E-commerce. In *Proceedings of ACM Conference on Electronic Commerce, Denver, Colorado, United States* (pp. 158–166).

Shardanand, U., & Maes, P. (1995). Social Information Filtering: Algorithms for Automating "Word of Mouth". In *Proceedings of International Conference on Human Factors in Computing Systems, Denver, Colorado, United States* (pp. 210–217).

Terveen, L. G., Hill, W. C., Amento, B., McDonald, D., & Creter, J. (1997). Building Task-Specific Interfaces to High Volume Conversational Data. In *Proceedings of Conference on Human Factors in Computing Systems, Atlanta, Georgia, United States* (pp. 226–233).

Tveit, A. (2001). Peer-to-peer Based Recommendations for Mobile Commerce. In *Proceedings of International Workshop on Mobile Commerce, Rome, Italy* (pp. 26–29).

Ungar, L. H., & Foster, D. P. (1998). Clustering Methods For Collaborative Filtering. In *Proceedings of the Workshop on Recommendation*

Systems.

Wang, J., Pouwelse, J., Lagendijk, R. L., & Reinders, M. J. T. (2006). Distributed Collaborative Filtering for Peer-to-Peer File Sharing Systems. In *Proceedings of ACM Symposium on Applied Computing, Dijon, France* (pp. 1026–1030).

Zhang, S., Wang, W., Ford, J., Makedon, F., & Pearlman, J. (2005). Using Singular Value Decomposition Approximation for Collaborative Filtering. In *Proceedings of the Seventh IEEE International Conference on E-Commerce Technology* (pp. 257–264).