Trust- and Location-Based Recommendations for Tourism

Annika Hinze and Qiu Quan

University of Waikato, New Zealand
a.hinze@cs.waikato.ac.nz, quan.qiu.yz@gmail.com

Abstract. Recommender systems in a travel guide suggest touristic sites a user may like. Typically, people are more willing to trust recommendations from people they know. We present a trust-based recommendation service for a mobile tourist guide that uses the notion of directly and indirectly trusted peers. The recommendations combine information about the peers ratings on sights, interpersonal trust and geographical constraints. We created two trust propagation models to spread trust information throughout the traveller peer group. Our prototype supports six trust-based, location-aware recommendation algorithms.

1 Introduction

Travellers using a mobile tourist guide like to receive recommendations about interesting sights. Typical recommender algorithms were developed for product recommendations in e-commerce, which do not consider location context. In addition, the classical techniques (collaborative filtering and content-based filtering) are not very transparent to the users so that they wonder: *why did I get this recommendation?* Long-term usage of recommenders is often related to the user's confidence and trust in the system. We found that people typically turn to travel agents for recommendations but they have the greatest confidence in recommendations by friends and others travellers they meet on a journey. We strongly believe that the information domain also plays an important role – data used for tourist recommendations has a location context and is thus different compared to commercial products. A tourist's preferences change depending on their location: e.g., interest in museums while visiting Paris and interest in outdoor activities in New Zealand. We identified five factors that should be taken into account: location, personal interest, travel history, feedback from peers, and trust in peers. This paper introduces recommendation algorithms that consider these factors.

Our goal is to increase the user loyalty, understanding and trust in a tourist recommender system. We therefore first analysed who users trust for travel recommendations. Our user study comprised a user study on three aspects: the typical sources of travel recommendations, the perceived reliability of recommendations based on their source, and the propagation of trust within a community. We performed four questionnaire-based surveys. Here we use selected results of the first survey focussing on perceived reliability of travel recommendation sources. The other surveys explored trust and similarity, trust in indirect friends, and design issues for trust-based recommenders. The questionnaires and a complete analysis is provided in [12].

R. Meersman, T. Dillon, P. Herrero (Eds.): OTM 2009, Part I, LNCS 5870, pp. 414–422, 2009.

The results of the first survey confirms that people prefer recommendations from friends. The most trustable friends are those who are most similar but people are also interested in recommendations from friends with different interests. With no recommendations from their friends, people are willing to consult tourist experts or geographically close travellers, but require a clear explanation of the results. Particularly interesting is that in addition to similar interests, also similar locations (current or in travel history) build trust. To support recommendations from tourist peers, we need to present users with transparent recommendation from geographically trusted users.

2 Trust-Based Recommendations in Tourism

2.1 Terms and Definitions

Following the results of our user study, we use the notions of *local trust* and *global trust* (=reputation) and additionally introduce the concepts of *geographic trust* and *location-aware reputation*. Two local trust metrics are *Personal trust* and *person-group confidence*. Personal trust refers to the subjective trust from one user to the other. Person-group confidence is a set of trust statements from a user to a direct peer a sight group (i.e., a set of sights) A *peer* p refers to a user. A *peer group* $P = \{p_1, p_2, \ldots, p_n\}$ represents a tourist community in the system. The corresponding peer group for a peer p_i ($p_i \in P$) is referred to as P_i, with $P_i \subseteq P$. Peers issue *personal trust* T_p to direct acquaintances: $T_p : P \times P \rightarrow [0, 1] \subset \mathbb{R}$. $T_p(p_a, p_b) = 1.0$ with $p_a, p_b \in P$ means the source peer p_a completely trusts the target peer p_b. If $T_p(p_a, p_b) = 0$, all information coming from this peer will will be ignored. Peers outside the peer group have a $null$ trust value. Note that personal trust is not symmetrical. Different to our approach, the Advogato [9] trust metric only allows *Boolean* trust; this directly classifies the peer groups into entrusted or distrusted ones. The reputation of peer p constitutes a global *reputational trust* T_r, it is calculated by averaging the received personal trust values from the other users: $T_r : P \rightarrow [0, 1] \subset \mathbb{R}$.

If the distance between the source peer and the target peer is less or equal to a given threshold λ, the target peer is given a *Geographic trust* T_g by the source peer. $T_g : P \times P \rightarrow [0, 1]$ with $T_{g_{(p_i, p_j)}} = 1 - \frac{distance(p_i, p_j)}{\lambda}$ and $i, j \in [1, n]; p_i, p_j \in P; distance(p_i, p_j) \leq \lambda$ where p_i is the source peer, and p_j is the target peer. Target peers who are closer to the source peer receive a higher geographic trust since their domain knowledge may benefit the source peer more. Users with high *location-aware reputation* T_{lr} have visited a given location the peer is interested in: $T_{lr} : P \times P \rightarrow [0, 1]$. The user-defined *Trust threshold* μ is used to identify individual trustable peers; this restricts the peer group when propagating trust on the trust network.

A *Trust path* graph ρ consists of two finite sets (see Fig. 1): a vertex set $V(\rho) \subseteq P$ for the peers and a directed edge set $E(\rho)$ representing the trust relationships [5]. The personal trust value is the weight of the edge issued by a to b. A trust path is a finite sequence of adjacent edges connected via vertices: $\rho[p_i, p_j] = p_i - p_k - \ldots - p_j$; $p_i, p_k, \ldots p_j \in P$ where p_i represents the source peer and p_j represents the target peer. All vertices (peers) on one path have to be unique.

Personal trust can only be propagated along the trust path if the peer has at least one (trusted) friend. Direct friends of the source peer are all adjacent peers on the trust paths which are connected by exactly one directed edge. The remaining peers on the same trust path are indirect peers of the source peer. The length of a trust path is calculated in the number of the edges between two peers; e.g., $\rho[p_a, p_e]$=$p_a - p_b - p_e$ is $s = 2$. As the trust path of the source peer to the target peer becomes longer, the trust decreases gradually. This trust decay D is defined as: $D : P \times P \rightarrow [0, 1] \subset \mathbb{R}$ The trust decay to each direct friend is 1. The value $d \in [0, 1]$ is the *decay constant*. Close friends have more influence on the recommendations. The community trust network graph N is formed by connecting direct trust paths of all personal trust network graphs N_i.

The *sight set* $S = \{s_1, s_2, \ldots, s_n\}$ contains all uniquely identifiable sights in the system. Sight group set $G = \{g_1, g_2, \ldots, g_k\}$ contains all classifications of sights. Each sight group $g_i \in G$ represents a category to which a sight $s_i \in S$ might belong. A nearby-sight group S_λ, $S_\lambda \subseteq S$, contains sights that are geographically close to the location of the user. λ is the distance threshold utilized for constraining nearby sights. The user feedback f is a numeric value between 0 and 10 regarding a sight. Trust-based recommendation R_{p_i} for a particular user p_i is an ordered list of the top N recommended nearby sights that have been associated with computed scores and a list of trustable peers: $R_{p_i} = \{(s_{\lambda_1}, f_1, P_{i_1}), \ldots, (s_{\lambda n}, f_n, P_{i_n})\}$, where $s_{\lambda_n} \in S_\lambda$ is one of the nearby sights, f_n is the computed score given to this sight. $P_{i_n} \subseteq P_i$ is the set of peers that recommended this sight.

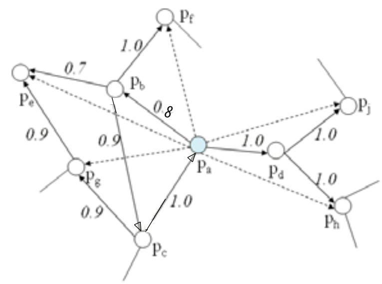

Fig. 1. Trust path graph

2.2 Trust Propagation Model

We designed three trust propagation models for the trust-based recommendation in TIP.

Model 1: Propagating Boolean trust. This trust model is the implementation of the most well-known local group trust metric, the Advogato metric [14]. Every user issues a Boolean decision (0 or 1) about the trustworthiness of direct peers. Trust will only be propagated if the personal trust of the peer is 1.

Model 2: Propagating numeric trust. Each user issues trust values to direct friends. Propagation of trust allows for predicting trust in indirect peers in two steps:

1. Calculation of trust on a given path. Three steps to calculate the trust value T between source and target on path ρ: (1) The propagated personal trust of source peer p_s to to target peer p_t is the product of all personal trust values on a path between them: $T_p(p_s, p_t) = T_p(p_s, p_i) * T_p(p_i, p_j) * \ldots * T_p(p_n, p_t)$, where $T_p(p_i, p_j)$ is the personal trust value issued by p_i to p_j. (2) The trust decay from the source peer to the target peer is computed from the number of edges on the path between two peers. (3) The trust value T from source peer to target peer is gained by multiplying the propagated personal trust T_p with the trust decay D. $T(p_s, p_t) = T_p(p_s, p_t) * D[p_s, p_t]$.

2. Trust path selection. If there are several paths existing between the source peer and the target peer, one needs to be selected to propagate the trust. By definition a trust path does not contain cycles. Trust path selection chooses the trust path with the overall maximum trust flow. First we need to identify all trust paths between source and target. Then the trust on each path is calculated. Finally, the path with the maximum trust value is selected: $T(p_i, p_j) = Max_{\forall\rho[p_i,p_j]}\{\prod_{k=1}^{w} T_p(p_i, p_j)(1-d)^{s-1}\}$, where w is the number of trust paths that connect the source peer and the target peer, $\forall\rho[p_i, p_j]$ includes all w paths between p_i and p_j, $p_i, p_j \in P$, and $p_j \in P_i$.

Model 3: Propagating trust person-group confidence. This is a modification of the previous model: Personal trust information as well as person-group confidence is propagated. This will generate more precise information of the user's preferences.

1. Person-group confidence definition. We identify the personal confidence regarding the sight groups ($\theta = \{g_1, g_2, \ldots, g_m\}$, the immediate sight categories of sights) of the friend as the *person-group confidence* vector.. $C_\theta^{p_i}(p_i, p_j) = \{c_{g_1}^{p_i}(p_i, p_j), \ldots, c_{g_m}^{p_i}(p_i, p_j)\}$ where $p_i, p_j \in P$, $p_j \in P_i$ and $g_m \in \theta$. Each element in the vector represents a confidence value on a particular sight group. Similar to personal trust, the person-group confidence needs to be explicitly issued to direct acquaintances: $C_\theta^P : P \times P \times \theta \rightarrow [0,1]$ User p_i may hold a person-group confidence matrix $C_{Matrix_{p_i}}$ of friends: $C_{Matrix_{p_i}} : P \times \theta \rightarrow [0,1]$. By using the person-group confidence, users can precisely express their preferences.

2. Propagating the person-group confidence on the trust network. The person-group confidence in an unknown peer is calculated similar to propagating trust: (1) select the trust path with the maximum trust flow between source peer p_1 to target peer p_n; (2) obtain person-group confidence vector issued by p_{n-1} to p_n; (3) update the person-group confidence vector using the maximum trust flow. The resulting vector is the computed *trust person-group confidence vector* of the source to the target. The trust person-group confidence is the combination of the person-group confidence and trust: $C_\theta : P \times P \times \theta \rightarrow [0,1]$ with $C_\theta(p_i, p_j) = C_\theta^{p_i}(p_i, p_j)T(p_i, p_j)$, $p_i, p_j \in P$.

2.3 Recommendation Generation

We present four base approaches of trust-based recommendations. The fourth approach uses three different options of how to mitigate missing feedback scores.

Approach 1: Boolean trust propagation. This approach involves three steps. The first step is trust propagation in which user p_1 can find a trust peer group P_1. The second step is to collect the history data H_{p_1} of feedback of the peer group regarding all nearby sights. In Step 3, recommendations are generated based on the peers' historic data only. The data in H_{p_1} is aggregated by using the average of the feedback issued by all friends for the same sight. The recommendation generation occurs after completing the data collection. In this approach, the recommended sights from direct and indirect recommenders are treated equally.

Approach 2: Numeric trust propagation. This approach integrates the level of trust into recommendation generation; it involves four steps. The first step is the trust propagation to identify the peer group of the user – similar to approach 1. The second step is to collect the historic data set H_{p_i} from all peers as described above. The third step is to integrate the trust into H_{p_i} to form a trust-based rating data set $H_{p_i}^t$. The trust matrix is used to update the data set H_{p_i} by multiplying the corresponding trust values with the ratings from peers. In step 4, the recommended sights are generated based on $H^t_{p_i}$. Recommendations are generated from most trustable peers.

Approach 3: Propagation of person-group confidence. Often a person may trust someone only in some aspects; the trust person-group confidence vector is used to express this. Four steps are needed for the recommendation process. Firstly, the peer group P_i is created for the user; secondly, the trust person-group confidence matrix $C_{Matrix_{p_i}}$ is computed; thirdly, the trust person-group confidence matrix is used to update the extracted historic data set from H_{p_i} to $H^c_{p_i}$; finally recommendations are generated from the confidence-matrix-based historic data set. The updated historic data set contains the information about user preferred recommenders and the user preferred recommended subjects; which better models the user's interests and more easily accepted by the user. The computational cost in the recommending process will be higher.

Approach 4: Trust and location-aware collaborative filtering. Pure collaborative filtering (CF) is based on the similarity between the source peer and the target peer. In our case, we assume that if A trusts B, a similarity between A and B exists. The recommendation of trust-enhanced collaborative filtering is a list of the top N appreciated near sights generated from most trustable and similar peers. This approach needs four steps to generate recommendations. Firstly the trust propagation; then the user similarity computation; thirdly the prediction of ratings for the active user; and in the final step the use of the trust matrix to weight predicted ratings of recommended sights. For the user similarity calculation (Step 2), the rating statements of each peer are taken from the current user's peer group P_i as input from the user similarity metric, and the output is the similarity value of current user p_i against the peer. We use the *Pearson Correlation coefficient*, which is between $+1$ and -1.

Handling of missing feedback scores. Since each user only rates some sights, the resulting rating matrix is sparse. We replace no-feedback values by either a high score or a neutral score based on three strategies using:

(F1) *information from the user profile:* A score of 10 is allocated to a sight that belongs to a sight group preferred by the user; a neutral score of 5 otherwise.

(F2) *information from travel history:* If the user has visited a sight at least twice, the no-feedback value replaced by a score of 10; or with 5 otherwise.

(F3) *information from both profile and travel history:* a disjunction of F1) and (F2).

We expect the computational cost to be much lower than pure CF. Moreover, users know the information source and are able to influence the process. Hence this solution is more transparent and controllable than pure collaborative filtering.

3 Evaluation

We implemented our 6 recommender strategies in the context of the TIP project [6]. We refer to the implementation of Approaches 1–3 as TL-1–3, respectively, and to the variations of Approach 4 as TCR-1–3. We examined the strategies in a quantitative

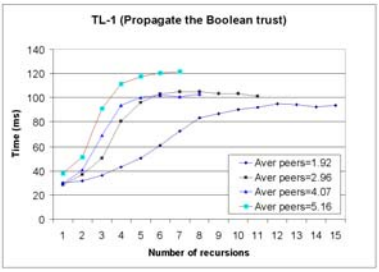

Fig. 2. Performance TL-1

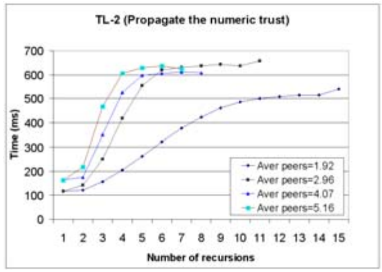

Fig. 3. Performance TL-2

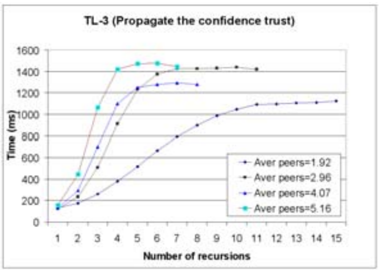

Fig. 4. Performance TL-3

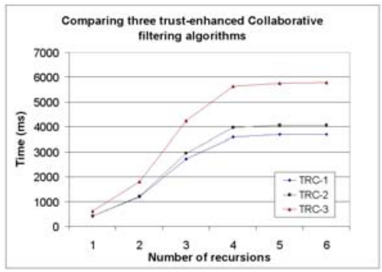

Fig. 5. TLC1–TLC3

performance evaluation as well as in a qualitative exploratory study on system functionalities to investigate the *transparency* of trust-based recommendations and *controllability* of the recommending process. Here we can present only a few results of our performance evaluation; for an extensive discussion see [12].

We varied the scale of the direct peer group to analyze time efficiency. All experiments were running on an Intel Pentium CPU 2.80GHz with 1.00GB RAM under Windows XP. In the test data sets, 100 users were created with each user having a direct peer group. The maximum number of peers in each direct peer group was 3, 5, 7 and 10. 500 unique sights were considered, equally classified into 10 sight groups. The maximum number of ratings issued by each user was 10. The direct trust value issued was a random real number value between 0 and 1, and the rating value was a random integer value between 1 and 10. Figs. 2 to 4 show that the response time for the three trust-based filtering algorithms TL-1 to 3 increases linearly with the number of peers in the peer group. The locations of recommended sights were not restricted. We see that performance depends heavily on the trust propagation model. The Boolean model (in TL-1) is a very simple and fast strategy. In TL-2, the trust model distributes the numeric trust on the network. This model needs to predict the trust of every indirectly trusted peer. In addition, trust values are integrated into peers' ratings from which recommendations are generated. As a result, the second model involves additional computation and is about five times slower. The trust model used in TL-3 specifies the trust on each particular sight group. In addition to propagating trust, more computation is needed for computing trust on different sight groups. The number of sight groups is an important factor that increases the time of trust propagation.

Fig. 5 shows the response times for the three trust-enhanced collaborative filtering algorithms TLC-1–3. The cost of computing user similarities and ratings prediction are similar as they use the same trust propagation model and sight rating matrixes ($users \times sights$). The main difference is the strategies for filling the no-feedback value. Response time for TLC-1 is influenced by the time to check the number of sight groups in the user's profile. TLC-2 needs to check the user's historic data, and then counts the number of times that each sight has been visited by the user. TLC-3 needs to check both the user profiles and user histories.

4 Related Work

Context-sensitive electronic tour guides. We analysed a number of tourist information systems [2,4,11,7,13] and found that none of them supports a rating component or collection of user feedback. Secondly, the displayed information follows a scheme that was pre-defined in the system or in users' profiles; other potentially interesting kinds of information would not be presented to users. Thirdly, users did not have access to information about other travellers' behaviours on similar travels nor to their comments.

Trust-based recommender systems. The trust concept has been evaluated to improve e-commerce recommenders, but most projects remain theoretical. The only running trust-based tourist recommender we encountered is the Moleskiing project [3] for skiing recommendations. Their trust metric MoleTrust [10] computes the user's trust in an unknown peer by tracing their personal social network. Users of Moleskiing need to manually record their skiing routes and experiences after finishing their trip. This prototype is not suitable for supporting ongoing travels. Ziegler and Lausen propose decentralized social filtering that uses trust network structures [14] for recommending books. The project combines the Appleseed trust propagation model (an adapted Advogado model) with a taxonomy-driven filtering technique to deal with data sparseness based on trust. The Appleseed model integrates the numeric trust weight, trust decay and trust normalization into the trust propagation, which make rankings feasible. Appleseed operates on a trust graph where a predefined trust threshold is used to detect entrusted peers in the process of exploring the social network. However, recommending sights is different to recommending books. In travelling, tourists' interests might change frequently according to the location, weather, season. Context-dependent changes of preferences have not been considered. The Recommend-Feedback-Re-recommend (RFR) conceptual framework [8] learns the users' preference from an FOAF-based environment. The system keeps updating distributed users' profiles to find a reliable user group. This framework provides a solution to group similar users in the distributed environment and recommendations automatically transfer to the other users starting from an initial recommender. Recommendations can only be transmitted to users through similar users who have rated the same item and where the given ratings satisfy the threshold. A distributed trust model [1] proposes a conditional transitive trust propagation model for a distributed environment and a protocol for recommendations: Trust is only allowed to be propagated if some trust conditions have been met.

5 Summary

Our initial research demonstrated that travellers prefer recommendations either from their friends or from people whose travel habits are similar to theirs. However, in practice these sources of information are not typically those most used. People are also willing to consult tourism experts or travellers nearby, but they then require a clear explanation of recommendations.

The goal of the research presented in this paper was to undertake a systematic interrogation of trust models in tourist recommender systems.

We proposed two trust concepts specifically designed for the tourist information domain: location-aware reputation and geographic trust. Geographic trust is higher the shorter the the distance between the active user and another user. Location-aware reputation is the reputation of users who have visited a certain location. These concepts give alternatives to users who have not defined their friends to the system or who do not agree with recommendations of their friends.

We created two trust propagation models and implemented six trust-based recommender algorithms. Our first trust propagation model uses direct numeric trust between users. In the second model a user's trust expresses confidence in users for certain sight groups. For comparison, we also implemented the well known Boolean trust model (Levien's Advogato local trust metric).

We evaluated three trust-based filtering algorithms and three trust-enhanced and location-aware collaborative filtering algorithms based on the second trust model. The efficiency of the trust-based filtering algorithms depends on the trust propagation model, while trust-enhanced collaborative filtering depends on the method to replace no-feedback values in the rating matrix. Due to its complexity, the trust-enhanced collaborative filtering is slower than the trust-based filtering.

The focus of this paper was local trust. In future work, we can compare the local trust metric against the global trust metric (reputation), and weigh trust-based recommendations constructed based on the local trust against recommendations created on reputation.

References

1. Abdul-Rahman, A., Hailes, S.: A distributed trust model. In: New Security Paradigms Workshop, pp. 48–60. ACM Press, New York (1998)
2. Abowd, G., Atkeson, C., Hong, J., Long, S., Kooper, R., Pinkerton, M.: Cy-berguide: A mobile context-aware tour guide. In: ACM Wireless Networks (1997)
3. Avesani, P., Massa, P., Tiella, R.: Atrust-enhanced recommender system application: Moleskiing. In: ACM Symposium on Applied Computing (2004)
4. Cheverst, K., Davies, N., Mitchell, K., Friday, A.: Experiences of developing and deploying a context-aware tourist guide: The guide project. In: Proc. of MOBICOM 2000 (2000)
5. Epp, S.: Discrete Mathematics With Applications. Brooks Publishing Company (1996)
6. Hinze, A., Voisard, A., Buchanan, G.: Tip: Personalizing information delivery in a tourist information system. Journal on Information Technology and Tourism 11(4) (2009)
7. Hristova, N., O'Hare, G.: Ad-me: A context-sensitive advertising system within a mobile tourist guide. In: Conf. on Artificial Intelligence and Cognitive Science, AICS (2002)

8. Kim, H., Jung, J.J., Jo, G.: Conceptual framwork for recommendation system based on distributed user ratings. In: Li, M., Sun, X.-H., Deng, Q.-n., Ni, J. (eds.) GCC 2003. LNCS, vol. 3032, pp. 115–122. Springer, Heidelberg (2004)
9. Levien, R., Aiken, A.: Attack-resistant trust metrics for public key certification. In: USENIX Security Symposium, pp. 229–242. USENIX Assoc. (1998)
10. Massa, P., Avesani, P.: Controversial users demand local trust metrics: An experimental study on epinions.com community. In: Nat. Conf. on Artificial Intelligence (2005)
11. O'Grady, M., O'Rafferty, R., O'Hare, G.: A tourist-centric mechanism for interacting with the environment. In: Proc. of MANSE 1999 (1999)
12. Qiu, Q., Hinze, A.: Trust-based recommendations in for mobile tourists in TIP. Technical report, University of Waikato (June 2007)
13. Youll, J., Morris, J., Krikorian, R., Maes, P.: Impulse: Location-based agent assistance. In: International Conference on Autonomous Agents (2000)
14. Ziegler, C., Lausen, G.: Paradigms for decentralized social filtering exploiting trust network structure. In: Meersman, R., Tari, Z. (eds.) OTM 2004. LNCS, vol. 3291, pp. 840–858. Springer, Heidelberg (2004)