

People-to-People Recommendation Using Multiple Compatible Subgroups

Yang Sok Kim, Ashesh Mahidadia, Paul Compton, Alfred Krzywicki,
Wayne Wobcke, Xiongcai Cai, and Michael Bain

School of Computer Science and Engineering
University of New South Wales
Sydney NSW 2052, Australia

{yskim, ashesh, compton, alfredk, wobcke, xcail, mike}@cse.unsw.edu.au

Abstract. People-to-people recommendation aims at suggesting suitable matches to people in a way that increases the likelihood of a positive interaction. This problem is more difficult than conventional item-to-people recommendation since the preferences of both parties need to be taken into account. Previously we proposed a profile-based recommendation method that first uses compatible subgroup rules to select a single best attribute value for each corresponding value of the user, then combines these attribute value pairs into a rule that determines the recommendations. Though this method produces a significant improvement in the probability of an interaction being successful, it has two significant limitations: (i) by considering only single matching attribute values the method ignores cases where different attribute values are closely related, missing potential candidates, and (ii) when ranking candidates for recommendation the method does not consider individual behaviour. This paper addresses these two issues, showing how multiple attributes can be used with compatible subgroup rules and individual reply rates used for ranking candidates. Our experimental results show that the new approach significantly improves the probability of an interaction being successful compared to our previous approach.

Keywords: recommender systems, social network analysis.

1 Introduction

Social networks connect people in the world. To extend them, many social network web sites have been developed. People in online social networks communicate using messages. To establish a successful social interaction, a user, called here a **sender**, sends a message to another user, called here a **receiver**, and the receiver should reply positively to the sender. The success of the interaction depends on **reciprocal interests** between the sender and the receiver. Therefore it is desirable for social networks to provide sophisticated services that encourage successful social interactions between users. People-to-people recommenders address exactly this issue. People-to-people recommenders differ from the conventional item-to-people recommenders since they need to take into account the preferences of both sender and receiver.

This research focuses on the problem of recommending suitable matches in an online dating social network. Unlike other social networks, online dating site members usually provide accurate and detailed information about themselves since this is the primary source of information potential partners use to find them. People who have (one or more of) the same attribute values are called here a **subgroup**, and it has been reported that people in the same subgroup usually have similar preference patterns [1]. Our people-to-people recommendation method, called **Compatible Subgroup Rules (CSR)** [2], is based on this property. The term ‘compatible subgroup’ means that a subgroup is more interested in another subgroup compared to other subgroups, and vice versa. The degree of compatibility is measured by a **compatibility score**, which is calculated using past interaction data (training data). A rule is defined to suggest a compatible subgroup for a given subgroup containing the active user. Recommendation rules are constructed from training data that contains information for senders, receivers, interaction results (positive or negative), and attribute values of the sender and the receiver, as follows. For a given subgroup defined by an attribute value of the active user, the method finds its compatible subgroup defined by a value of the *same* attribute using the training data. The attribute values of the identified compatible subgroups are incrementally combined together to increase the probability of success interaction. This attribute combination process finishes if there is no significant interaction between a user subgroup and its compatible subgroup in the training data. The most specialised compatible subgroups are used to propose recommendations. This method shows a 26% success rate for the top 10 recommendations and 23% for the top 100 recommendations, considerably greater than the 15% baseline success rate [3].

In spite of its success, this previous method has the following limitations. First, the method only chooses a *single* compatible subgroup for each subgroup defined by the active user’s attribute value. It ignores the fact that an attribute may have multiple values that result in very similar compatibility scores to the active user subgroup. Since candidates having the similar (but not best) attribute value are not recommended, many suitable candidates for recommendation may be missed. Second, the method only measures a compatibility score using the sender’s interest in the receiver subgroup without considering other possible interest measures. Finally, the method ranks the candidate receivers using the compatibility scores of the *age* and *location* attributes, which does not take into account individual interactions. This paper addresses these limitations and presents several extended CSR methods. The extended CSR methods utilise multiple compatible subgroups based on heuristics for including multiple values for an attribute, include various compatibility score measures, and consider the individual activities of compatible subgroup members to rank the recommended candidates. We conducted experiments to evaluate the extended CSR methods using historical data from a commercial online dating site. Our experimental results show that the best new proposed method significantly improves the probability of successful interactions compared to the CSR method based on single attribute values.

2 Related Work

A recommender system is a system that learns a user’s preferences and makes recommendations the user might like. Recommender systems can be classified into

item-to-people and people-to-people recommenders [4,5]. Item-to-people recommendation has been studied from the early stages of the Web, and has been applied in various domains. Although there is some non-Web research [6,7], people-to-people recommendation has recently gained attention as social networks become popular on the Web [8–10]. People-to-people recommender systems have been used to recommend experts [7,11,12], friends [10,13] and collaborators [14], etc. Our research has focused on recommending partners in online dating. Methods for people-to-people recommendation can be classified into behaviour-based methods, profile-based methods and hybrid methods. Behaviour-based methods exploit interaction data between users such as contact between users, comments posted on another user’s page, following relationships to other users, etc. Typically a collaborative filtering approach is used to identify similar users and generate candidate users using interaction data [15,16]. Profile-based methods use user-supplied data to suggest candidate recommendations [6,11–13,17]. Hybrid methods combine behaviour-based and profile-based methods [8]. We previously proposed a profile-based recommendation method that uses subgroup interaction data to find a compatible subgroup based on a user’s profile [2]. We also proposed hybrid methods that combine our interaction-based collaborative filtering approach with our profile-based method [3]. In this paper, our research focuses on extending our previous profile-based approach, but it could easily be combined with interaction-based collaborative filtering in ways similar to those described in [3].

3 Recommendation Methods

3.1 Definitions

Definition 1. A **user** u is a member of the social network website and is described by n distinct attribute values for n different attributes:

$$u = \{u_1, \dots, u_n\}$$

where u_i represents the value of attribute i . The user provides these attribute values when joining the site and can update them at any time. The site allows the user to select only one value for each attribute. Regardless of type, e.g. numeric for number of children or date for date of birth, all values are converted into nominal values before recommendation rules are constructed.

Definition 2. A **subgroup** is a group of users who have the same gender and the same values for one or more attributes. For example, the female postgraduate subgroup consists of the female users who have value *postgraduate* for the *education* attribute.

Definition 3. A **contact** is an action where a user, the **sender**, sends a message to another user, the **receiver**. Such a contact indicates that the sender is interested in the receiver.

Definition 4. An **interaction** between two users is a contact from a sender to a receiver with either a reply from the receiver to the sender or no reply (if the receiver chooses to ignore the contact). An interaction is classified as either positive or

negative depending on the receiver's reply message type (which is defined by the system), and if the receiver does not reply to the sender, the interaction is classified as negative. An interaction always shows that the sender is interested in the receiver, but the receiver may or may not be interested in the sender.

Definition 5. A **sender (receiver) subgroup** is the set of senders (receivers) in a set of interactions. Sender subgroups are denoted sg ; receiver subgroups are denoted rg .

3.2 Interest Measures

Interest measures show how much the members of one subgroup are interested in the members of another, and are calculated using the number of interactions between them. Interest measures in this work are computed from interaction data for the training period (three months before the time of recommendation). We define four types of interest measure reflecting different intuitions.

Definition 6. Interest in target group measures how much more a sender subgroup sg contacts a receiver subgroup rg compared to all receivers R . This is defined as:

$$P(sg, rg) = \frac{n(sg, rg)}{n(sg, R)}$$

where $n(sg, rg)$ is the number of messages sent from the sender subgroup sg to the receiver subgroup rg and $n(sg, R)$ is the number of messages sent from the sender subgroup sg to all receivers R .

Definition 7. Lift in interest in target group measures how much more a sender subgroup sg , compared to all senders S , is interested in a receiver subgroup rg . This is defined as:

$$L(sg, rg) = \frac{P(sg, rg)}{P(S, rg)}$$

where $P(sg, rg)$ is the sender subgroup sg 's interest in the receiver subgroup rg , and $P(S, rg)$ is all sender S 's interest in the receiver subgroup rg , as in Definition 6.

Definition 8. Success with target group measures how much more successful a sender subgroup sg is when contacting a receiver subgroup rg compared to all receivers R . This is defined as:

$$P(sg, rg, +) = \frac{n(sg, rg, +)}{n(sg, rg)}$$

where $n(sg, rg, +)$ is the number of messages sent from the sender subgroup sg to the receiver subgroup rg with positive replies from the receivers, and $n(sg, rg)$ is the number of messages sent from the sender subgroup sg to the receiver subgroup rg .

Definition 9. Lift in success with target group measures how much more successful a sender subgroup sg is when contacting a receiver subgroup rg compared to all receivers R , relative to the improvement in success for all senders S . This is defined as:

$$L(sg, rg, +) = \frac{P(sg, rg, +)}{P(S, rg, +)}$$

where $P(sg, rg, +)$ is the sender subgroup sg 's success with the receiver subgroup rg and $P(S, rg, +)$ is the success of all senders S with the receiver subgroup rg .

3.3 Compatibility Score

In this research, the term 'compatibility' means that a sender subgroup is interested in a receiver subgroup, and vice versa. The compatibility score measures the degree of compatibility of two subgroups. To dampen significant differences between two subgroup interests in each other, we use the harmonic mean to combine the two reciprocal interest measures.

Definition 10. A **compatibility score** between a sender subgroup sg and a receiver subgroup rg measures how much two subgroups are interested in each other. It is defined as:

$$cs(sg, rg) = \frac{2I_{sg,rg}I_{rg,sg}}{I_{sg,rg} + I_{rg,sg}}$$

where $I_{sg,rg}$ is the interest of sender subgroup sg in the receiver subgroup rg and $I_{rg,sg}$ is the interest of rg in sg . Note that there are different compatibility scores for a sender and receiver subgroup depending on the particular interest measure.

3.4 Learning Compatible Subgroup Rules

Compatible subgroup rules are used to recommend candidate receivers for a given user [2,3]. A compatible subgroup rule has the form:

if $a_1 = u_1$ **and** ... **and** $a_m = u_m$ **then** $a_1 = v_1$ **and** ... **and** $a_m = v_m$

where each a_i is an attribute (e.g. *occupation*) and each v_i is the best matching value of this attribute for candidates to the value u_i of the user. The meaning of the rule is that if a user satisfies the condition part of the rule, the system should recommend candidates satisfying the conclusion part of the rule.

The above approach is, however, limited in that for a given user attribute value, only a **single attribute value** of the candidate is chosen. When the candidate attribute values are closely related, many acceptable candidates are therefore not recommended. As an example, the *occupation* attribute has values *legal*, *accounting* and *consulting*, which the data indicates are all very similar in compatibility for the value of *legal* for the *occupation* attribute of the user. To allow for rules with multiple attribute values, we generalize the form of the compatible subgroup rule as follows:

if $a_1 = u_1$ **and** ... **and** $a_m = u_m$ **then** $a_1 \in V_1$ **and** ... **and** $a_m \in V_m$

where each a_i is an attribute (e.g. *occupation*) and each V_i is a set of compatible **multiple attribute values** of this attribute for candidates to the value u_i of the user (equivalently, we treat each condition $a_i \in V_i$ as a disjunction of equality tests).

Our recommendation algorithm based on compatible subgroup rules is as follows. Let a user have attribute values $\{u_1, \dots, u_n\}$ for attributes $\{a_1, \dots, a_n\}$. Rules are constructed for each user individually. Rule construction starts with an empty rule (no condition and no conclusion) and produces a sequence of more and more specific rules $R_1, \dots, R_m$. At each step, starting with a rule R , an attribute a with value u of the user is chosen that has the highest subgroup success rate; $a = u$ is added to the condition of the rule and $a \in V$ is added to the conclusion of the rule, where V is the set of compatible values for the gender (male/female) of the sender (see below). These conditions are added to the rule only if it improves the overall rule success rate $SR(sg, rg)$ defined as follows, and this improvement is statistically significant:

$$SR(sg, rg) = \frac{n(sg, rg, +)}{n(sg, rg)}$$

where sg is the sender subgroup defined by the current rule condition and rg is the receiver subgroup defined by the current rule conclusion.

The two main aspects of the recommendation algorithm are finding multiple compatible subgroups (making use of heuristics to combine single attribute values into sets of compatible attribute values), and generating and ranking candidate receivers.

Finding Multiple Compatible Subgroups. For a user with value u of attribute a , the set of compatible attribute values V is determined as follows. First, for all possible values v_i of a , the compatibility score of v_i with u is calculated using Definition 10, giving a distribution of scores. The mean and standard deviation of this set is calculated. A value is included in the set V if its score is at least 0.5 standard deviations above the mean (this heuristic is based on preliminary experimental analysis).

For example, Figure 1 illustrates compatibility scores of different female age subgroups with the male age 35–39 subgroup. The female age 35–39 subgroup is the most compatible subgroup. However, it also shows that the female age 30–34 subgroup has a very similar compatibility score to the age 35–39 subgroup. Both are included in the set V .




Fig. 1. Multiple Compatible Subgroups: Compatibility scores for male age 35–39 subgroup with female age subgroups. The threshold line shows the mean plus 0.5 standard deviations.

Attribute	Condition (User Value)	Conclusion (Compatible Values)	Rule Success Rate
gender	male	female	14.9%
age	42	35–39 OR 45–49 OR 30–34	14.7%
location	Sydney-North	Sydney-East OR Sydney-North	15.2%
education	Graduate	Graduate OR Diploma	32.7%
smoke	Don't smoke	Don't smoke	37.5%
...	...	...	...
drink	Occasionally	Occasionally	54.5%

Fig. 2. Compatible Subgroup Rules: Each row represents a pair of conditions successively added to the rule where multiple attribute values for each user attribute are chosen. As conditions are added, the rule success rate increases but the number of potential candidates decreases.

Figure 2 illustrates an example rule of rule construction for a hypothetical male user, starting with *gender*. First *age* is considered. Since the user is 42, he is in the age band 40–44; the best set of female age bands are 35–39, 45–49 and 30–34. Accordingly the disjunction of these conditions for the female is added to the rule. Rule construction continues with the *location* attribute, where the male is in Sydney-North. The set of compatible values for the female are Sydney-East and Sydney-North. Continuing this process to define a sequence of rules, the rules become more and more specific, but generate fewer and fewer candidates.

Generating and Ranking Candidate Receivers. Given the sequence of rules $R_1, \dots, R_m$ generated as above, the potential recommendations are all candidates satisfying the conclusion part of any such rule who were active in the previous 28 days (but only as many candidates as needed are generated, starting with the most specific rule R_m). Candidates satisfying a more specific rule R_i are ranked higher than those only satisfying a rule R_j where $j < i$. However, candidates satisfying the same rule are ranked as follows, using data from the previous 28 days. First, candidates are ordered in terms of positive reply rate (defined as the proportion of contacts received to which the candidate has given a positive reply). Second, if positive reply rates are equal, the number of positive replies is used to order the candidates (those with a larger number are ranked higher). Third, if two candidates are still equally ranked, they are ranked in order of the number of contacts received (again, those who received more contacts are ranked higher). Finally, any ties between remaining candidates are broken randomly.

4 Experimental Design

To evaluate the proposed new methods for defining compatible subgroup rules using multiple attributes with different interest measures and the new way of ranking candidates, we used the same experimental setup previously used for single attribute compatible subgroup rules [3]. In these experiments we used data from a commercial online dating site and generated recommendations once, for March 1, 2010, using only data that would have been available on that date. For evaluation of the recommendation methods, we collected two types of data – interaction data and user profile

data. Profile data consists of basic information about the user, such as age, location, marital and family status, plus attributes such as smoking and drinking habits, education and occupation. Each interaction is recorded with a date/time stamp, the type of message and the response message type, which is labelled as either positive or negative, or null if no reply has been received (contacts without reply are classified as negative). We used historical data to evaluate the recommendation methods. The following datasets were collected for the experiment:

- **Rule Learning Data:** This dataset consists of interaction data and user profile data for the users who were the senders and/or receivers of interactions for three months prior to the recommendation construction date. There were about 5 million interactions by 200,000 users. The collected data was used to learn the recommendation rules.
- **Active User Data:** This dataset consists of active user profiles and their interaction data for 28 days prior to the recommendation construction date. The collected dataset is used to generate candidate receivers using the recommendation rules. There are about 137,000 active users.
- **Test Data:** This dataset consists of interactions from the first three weeks of March 2010 initiated by the active users. The test set consists of around 130,000 users with around 650,000 interactions, of which roughly 15% are positive interactions.

After collecting these datasets, we generated up to 100 candidate receivers for each active user using our recommendation methods, with candidate receivers ranked as defined above. We then examined whether or not the data showed positive or negative test interactions with the candidate receivers. Note that we did not provide users with actual recommendations; the method compares the recommendations with the interactions that occurred in the test set with people who would have been recommended.

The two main metrics are precision and recall for the top N recommendations, where $N = 10, 20, \dots, 100$. **Precision** measures the proportion of recommended interactions that occurred which were positive, so is analogous to the user's success rate when adopting the recommendations. More precisely, if R is the set of recommended interactions and T is the test set, precision is defined as follows, where T^+ is the set of positive interactions in T and $n(S)$ the number of elements in the set S .

$$precision = \frac{n(R \cap T^+)}{n(R \cap T)}$$

A high precision suggests that the recommendations are useful, while a low precision suggests the recommendations are not helpful. Note that this measure can be considered the probability that an interaction resulting from a recommendation is positive. Therefore, we can compare this result with the baseline probability of an interaction being positive (without recommendation).

Recall measures how many of the positive test interactions are included within the top N recommendations, so is analogous to how much the users liked the recommendations. This is measured by the proportion of positive interactions recommended that are positive interactions in the test set.

$$recall = \frac{n(R \cap T^+)}{n(T^+)}$$

This measure indicates whether positive interactions are likely to happen when the recommendations are given to users. If the probability of an interaction being positive (precision) is very high, but such interactions are unlikely to occur, those recommendations are not particularly useful. Therefore, for recommendation methods where the precision is similar, a high recall is a good indicator that the recommendations are better. Even though we have used precision and recall as our evaluation metrics, they differ from the conventional use. Normally such measures are applied to a test set where each item is known to be relevant or irrelevant, but in our experiments we calculate these measures based on the interactions that occurred in historical data. Since these actually occurring interactions are only a very small subset of all possible interactions (around 650,000 out of 4.7 billion), and since we generate only a limited number of recommendations (only up to 100 for each user), the intersection of the recommended interactions and the test set interactions is very small. This results in very small recall results in our experiments. Nevertheless, recall is a useful metric for comparing different methods on how much users like the given recommendations.


5 Results

5.1 Precision and Recall with Candidate Behaviour-Based Ranking


The basic comparison of our methods is shown in Figure 3, which shows precision and recall for the top N recommendations for our previous method based on **single attribute values** [2,3] (ranked using the harmonic mean of the compatibility scores for *age* and *location* under Definition 6), and four new methods based on compatible subgroup rules using **multiple attribute values** with the four different interest measures (Definitions 6–9) and behaviour-based ranking as described in Section 3.4.

The first observation is that all methods based on multiple attributes have a much higher precision (for all values of N) than the previous method based on single attributes. Second, though the methods using interest measures based on successful interactions show the highest precision, recall is very low for these methods, suggesting that the rules used to generate the recommendations are too specific. In other words, the fraction of the candidate receivers that show a high success rate is usually very small, and combining several of those compatible values results in very successful interactions, but these interactions very rarely occur in the historical data.

Thus the methods based on just **multiple attribute values** with **interest in target group** and **multiple attribute values** with **lift in interest in target group** perform best, with similar precision and much higher recall than the methods using interest measures based on successful interactions. The method using **interest in target group** is slightly superior. Moreover, the recall is higher than that based on **single attribute values**, giving evidence that these methods based on **multiple attribute values** with behaviour-based ranking are generating better candidates.



(a) Precision




(b) Recall


Fig. 3. Precision and Recall with Candidate Behaviour-Based Ranking

5.2 Precision and Recall without Candidate Behaviour-Based Ranking

To examine whether the improvement over the previous single attribute method is due to the ranking or to the method of constructing rules by combining multiple attributes, precision and recall for the methods without using behaviour-based ranking are shown in Figure 4. In this case, the candidate receivers are only ranked by the compatible subgroup rules, with ties broken randomly (hence it is expected that precision will be lower). In fact, the precision for the methods based on **multiple attribute values** without success is much lower than in Figure 3, similar to that for **single attribute values**, though slightly higher with **lift in interest in target group**. Recall, however, is comparable to the method based on **single attribute values**. The results show that the ranking method is critical to the increased precision of the methods.



(a) Precision



(b) Recall

Fig. 4. Precision and Recall without Candidate Behaviour-Based Ranking

6 Conclusion

This paper presented several extensions of a profile-based method for people-to-people recommendation based on compatible subgroup rules. We proposed four types of interest measures, and a new way of combining several compatible attribute values in rule construction, and a ranking method based on candidate behaviour. Our experimental results show that a new method based on **multiple attribute values** with **interest in target group** has much higher precision than the other methods and higher recall than our previous method based on **single attribute values**, indicating that the recommendations generated using this method are of higher quality. Our results also show that candidate behaviour-based ranking is very effective in reordering the candidate receivers. In other work [3], we combined single attribute compatible subgroup rule recommendation with interaction-based collaborative filtering [15] and obtained significant improvements in precision and recall. We expect that our new methods are also likely to improve the performance of such hybrid recommendation methods, but this has not yet been evaluated. This will be the subject of further work.

Acknowledgements. This work was funded by Smart Services Cooperative Research Centre. We would also like to thank our industry partners for providing the datasets.

References

- [1] McPherson, M., Smith-Lovin, L., Cook, J.M.: Birds of a Feather: Homophily in Social Networks. *Annual Review of Sociology* 27, 415–444 (2001)
- [2] Kim, Y.S., Mahidadia, A., Compton, P., Cai, X., Bain, M., Krzywicki, A., Wobcke, W.: People Recommendation Based on Aggregated Bidirectional Intentions in Social Network Site. In: Kang, B.-H., Richards, D. (eds.) PKAW 2010. LNCS, vol. 6232, pp. 247–260. Springer, Heidelberg (2010)
- [3] Kim, Y.S., Krzywicki, A., Wobcke, W., Mahidadia, A., Compton, P., Cai, X., Bain, M.: Hybrid Techniques to Address Cold Start Problems for People to People Recommendation in Social Networks. In: Anthony, P., Ishizuka, M., Lukose, D. (eds.) PRICAI 2012. LNCS, vol. 7458, pp. 206–217. Springer, Heidelberg (2012)
- [4] Pizzato, L., Rej, T., Chung, T., Koprinska, I., Kay, J.: RECON: A Reciprocal Recommender for Online Dating. In: *Proceedings of the Fourth ACM Conference on Recommender Systems*, pp. 207–214 (2010)
- [5] Pizzato, L., Rej, T., Akehurst, J., Koprinska, I., Yacef, K., Kay, J.: Recommending People to People: The Nature of Reciprocal Recommenders with a Case Study in Online Dating. *User Modeling and User-Adapted Interaction* (2012)
- [6] Belkin, N.J.: Helping People Find What They Don't Know. *Communications of the ACM* 43(8), 58–61 (2000)
- [7] McDonald, D.W., Ackerman, M.S.: Expertise Recommender: A Flexible Recommendation System and Architecture. In: *Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work*, pp. 231–240 (2000)
- [8] Kazienko, P., Musiał, K.: Recommendation Framework for Online Social Networks. In: Last, M., Szczepaniak, P., Volkovich, Z., Kandel, A. (eds.) *Advances in Web Intelligence and Data Mining*, pp. 111–120. Springer, Berlin (2006)

- [9] Yu, S.J.: The Dynamic Competitive Recommendation Algorithm in Social Network Services. *Information Sciences* 187, 1–14 (2012)
- [10] Naruchitparames, J., Giine, M.H., Louis, S.J.: Friend Recommendations in Social Networks Using Genetic Algorithms and Network Topology. In: *Proceedings of the 2011 IEEE Congress on Evolutionary Computation (CEC)*, pp. 2207–2214 (2011)
- [11] Moraes, A., Silva, E., da Trindade, C., Barbosa, Y., Meira, S.: Recommending Experts Using Communication History. In: *Proceedings of the 2nd International Workshop on Recommendation Systems for Software Engineering*, pp. 41–45 (2010)
- [12] Shami, N.S., Yuan, Y.C., Cosley, D., Xia, L., Gay, G.: That's What Friends are For: Facilitating 'Who Knows What' Across Group Boundaries. In: *Proceedings of the 2007 International ACM Conference on Supporting Group Work*, pp. 379–382 (2007)
- [13] Chen, J., Geyer, W., Dugan, C., Muller, M., Guy, I.: Make New Friends, But Keep the Old: Recommending People on Social Networking Sites. In: *Proceedings of the 27th International Conference on Human Factors in Computing Systems*, pp. 201–210 (2009)
- [14] McDonald, D.W.: Recommending Collaboration with Social Networks: A Comparative Evaluation. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 593–600 (2003)
- [15] Krzywicki, A., Wobcke, W., Cai, X., Mahidadia, A., Bain, M., Compton, P., Kim, Y.S.: Interaction-Based Collaborative Filtering Methods for Recommendation in Online Dating. In: Chen, L., Triantafillou, P., Suel, T. (eds.) *WISE 2010. LNCS*, vol. 6488, pp. 342–356. Springer, Heidelberg (2010)
- [16] Pizzato, L.A., Silvestrini, C.: Stochastic Matching and Collaborative Filtering to Recommend People to People. In: *Proceedings of the Fifth ACM Conference on Recommender Systems*, pp. 341–344 (2011)
- [17] Cai, X., Bain, M., Krzywicki, A., Wobcke, W., Kim, Y.S., Compton, P., Mahidadia, A.: Learning to Make Social Recommendations: A Model-Based Approach. In: Tang, J., King, I., Chen, L., Wang, J. (eds.) *ADMA 2011, Part II. LNCS*, vol. 7121, pp. 124–137. Springer, Heidelberg (2011)