SRec: A Social Behaviour Based Recommender for Online Communities

Surya Nepal, Cecile Paris
CSIRO ICT Centre
PO Box 76, Epping NSW, 1710 Australia
{Firstname.Lastname}@csiro.au

Sanat Bista
CSIRO ICT Centre
Canberra, Australia
Firstname.Lastname@csiro.au

ABSTRACT
Recommender systems have been successfully used in electronic commerce applications such as recommending books, movies, restaurants and airlines based on users' past behaviour. More recently, such systems have made inroads into social media, for examples to recommend partners in online dating sites. In our work, we have developed a social behaviour based recommender system within an online community with the aim to increase the level of interactions in the community, thereby increasing its *social capital* (the density of interactions among its members in the community) and its chance of sustainability. Our recommender system is built on a social trust model. It is able to recommend people and content. Importantly, it can recommend people in different roles: friends, mentors and leaders. In this paper, we describe our context and the social behaviour based recommender system we developed.

Author Keywords
Social Trust, Social Behaviour, Recommender, Online Communities.

ACM Classification Keywords
H3.3 [Information Search and Retrieval]: Information filtering, Selection process. H.3.4 [Information Storage and Retrieval]: Systems and Software. H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous. H J.4 [Computer Applications] Social and Behavioral Sciences.

INTRODUCTION
Recommender systems have been successfully used in electronic commerce applications such as recommending books, movies, restaurants and airlines [1]. More recently, such systems have been used to recommend people for social interactions, for example to recommend partners in online dating sites [2], to recommend professional relationships in sites like Linkedin [3] or to recommend strangers in the enterprise [4].

Recommendation systems can be broadly categorised as (i) content based recommendation [6], (ii) collaborative filtering based recommendation [7], (iii) trust based recommendation [8] and (iv) hybrid recommendation systems [5]. Content-based approaches produce recommendations based on the similarity between items consumed by the members. Collaborative filtering approaches recommend the items chosen by the users with similar tastes/preferences. Trust based recommendation systems usually construct a trust network where nodes are users and edges represent trust between two users. The goal of a trust based recommendation system is to generate personalised recommendations by aggregating the opinions of users in their trust network. Hybrid systems support a combination of these three approaches. Our focus is on a trust based recommender system in the context of an online community in which we want the system to be able to recommend both people and resources to the community members. The aim of the recommender system is to increase the level of interactions in the community, thereby increasing the *social capital* (the density of interactions among its members in the community) of the community and its chance of sustainability. The recommender system also has two intents: reinforcement of (good) user behaviour and motivation for change (of attitude, for example).

Early trust based recommender systems [9] borrow the concept of trust from electronic commerce applications and peer-to-peer systems. A trust network is built using explicit indications of trust from one user to another, and trust can be propagated through the network via friend-of-a-friend (FOFA) relationships [10]. Later models have started to look at the social aspects of trust [8]. In online communities, we believe that the social aspect of trust, as defined in social and behavioural sciences [11], must be taken into account. This is what we are attempting to do in our work. To this end, we have developed a *social trust model*, *STrust* [12], itself inspired by notions of trust from Social Sciences. The *Strust* model derives trust from the social capital of a community as a whole as well as individuals' social trust. Our trust model supports a recommendation system, *SRec*, which is built on the social

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
3rd Workshop on Social Recommender System at UMAP, Montreal, Canada, July 16-20, 2012
Copyright 2012 XXX XXX-X-XXXXX-XXX-X/XX/XX...$5.00.

trust model and is capable of recommending both people and resources using contextual information. Importantly, *SRec* is able to recommend different types of people.

The rest of this paper is structured as follows. We first introduce the context of our work, the online community for which we have developed our social trust based recommendation system. We then briefly discuss our social trust model, focusing on the features on which our recommendation system rely. We then describe key components of our social recommender system (SRec). Finally, we present the state of our work and potential future directions.

OUR CONTEXT: AN ONLINE COMMUNITY FOR THE DELIVERY OF GOVERNMENT SERVICES

We are building an online community to deliver government services to citizens [13, 14], as a trial for this type of service delivery.

The key features of the community are:

- The community is by invitation only, i.e., only specific individuals (individuals receiving a specific type of welfare payments) are invited to join the community;
- People in the community do not know each other, and we do not know whom they are as the invitation process is double blind (following the rules of our ethics committee). For privacy reasons, members are discouraged to reveal their real identity. Members have profiles, but they disclose to others only what they wish to disclose.

Our specific target group is in a transition phase, being asked (by legislation) to move from one type of payment to another. As the result of the transition, the individuals involved receive less money and are required to go back to the workplace. Understandably, this is a difficult transition, and some individuals have a resentful attitude towards the government and their required return to work.

The aim of the community is several fold. First, it is to bring people with the same concerns together, hoping that they will share experiences, ideas and tips, thus providing social, emotional and moral support to each other. All individuals are stranger to each other – but they all share the same situation. Second, it is a space in which we invite individuals to go on a reflection journey, in order to prepare them better for the transition and their return to work. Finally, it is a place for the government to target its information and services when dealing with a specific target group of welfare recipients.

To encourage interactions in the community, we have developed a number of mechanisms, including activities, specifically designed to encourage reflection, a discussion forum, a set of information resources, a "buddy program" (or "Friends Circle", a content-based recommendation system to encourage people to find someone with whom they can do this journey), some gamification elements [15], and a Live Chat feature to enable individuals to converse in real time with an expert on a specific topic (e.g., how to prepare for an interview). It is also in this context that we have developed our social trust model, as we want to investigate how we can increase social trust amongst the individuals in the community, thereby increasing the community's social capital.

OUR SOCIAL TRUST MODEL: THE STRUST MODEL

We define social trust as "*the positive behaviour/interactions of users in the community at any particular time in a certain context*". This definition refers to three important aspects of trust.

User behaviour: Social trust depends on the behaviour of an individual, as derived from his or her interactions in the community through different types of activities such as writing, reading or commenting on a post, viewing information, participating in an activity, etc. We consider two types of interactions: *active* (e.g., befriending someone, writing a post, etc.) and *passive* (e.g., regular visits to the community, reading a post or an informational resource, etc.). Taken together, these two types of interactions collectively build the social capital of the community.

Temporal factor: Time plays an important part in social trust: a recent interaction has more weight than one that occurred a long time ago.

Context: Context for an interaction in our community refers to whether the interaction was performed in the discussion forum (e.g., reading a post), the Live Chat (e.g., someone participated in it), the resources (someone read or rated a resource), etc. All interactions are done in a specific context.

Our social trust model has three unique features. First, it distinguishes between three trust types: the popularity trust (*PopTrust*), that captures the trust that an individual member has received from other members in the community, the engagement trust (*EngTrust*), which reflects the trust that an individual member has about other members in the community, and the social trust (*Strust*), that combines the two.

The second feature of the *STrust* model is that it considers both active and passive behaviour of members. Active behaviour refers to actions that generate: (a) content for other members in the community to consume (e.g., contributions to forum, Live Chat, etc.), and (b) actions that require other members to act (e.g., invitation to be a friend). Passive behaviour refers to actions that generate neither content nor actions for other members (e.g., visiting the community, reading posts, reading Live Chat content, etc.).

The third feature of the *STrust* model is that it considers online communities as two mode social networks, where nodes in the networks can be classified into two types: *active* and *passive*. Active nodes are those that can engage

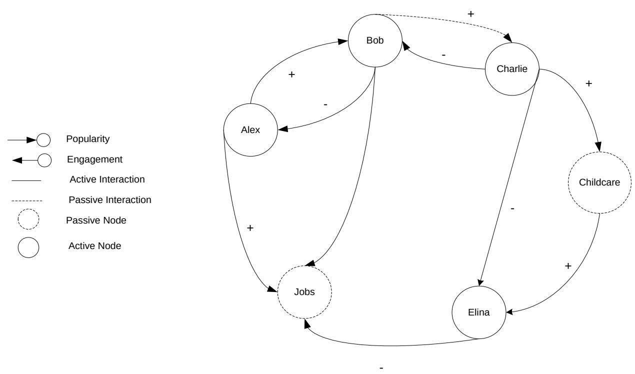

(a)

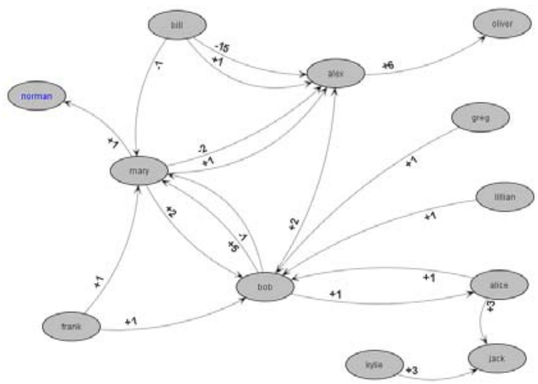

(b)

Figure 1. (a) Community Behaviour Model (b) Behaviour from our Online Community Development Server

in the community (typically, people), while passive nodes are those that cannot (i.e., they do not have engagement trust), such as articles and posts.

Figure 1 (a) shows a community behaviour model, capturing the different entities and behaviours just discussed. The nodes in the graph represent community members and the edges their interactions. The nodes could also be other entities in the community, e.g., tasks or contents, which we refer to as passive nodes. Each arrow provides information towards popularity trust for one side (the sink or receiving end) and engagement trust on the other side (the source or initiating end). Let's consider the node representing Bob as an example. It has three outgoing arrows and two incoming ones. The outgoing arrows support Bob's engagement trust, and the incoming ones support Bob's popularity trust. Solid lines represent active interactions, and dotted lines represent passive interactions. Finally, the interactions between two nodes are either positive (represented as +) or negative (represented as -). Passive interactions are always considered positive. Figure 1(b) shows the visualisation of the behaviours of test users in our community for two contexts: rating and comments.

The exact description of our STrust model is out of the scope of this paper. See [12] for details.

SREC: RECOMMENDER SYSTEM

Our recommender system, *SRec*, is a social behaviour based recommender system built upon the social trust model. As previously mentioned, its main purpose is to increase the social capital, that is, the density of the interactions in the community: it does so by recommending new activities that lead to a new set of interactions amongst community members and with the available resources.

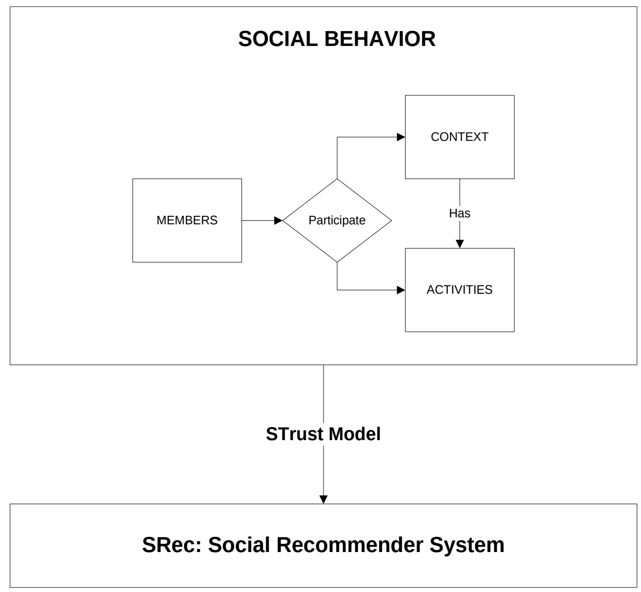

Figure 2: SRec – A social behaviour based recommender system

Figure 2 shows the social behaviour based recommender system. The first step is to capture the social behaviour of the members in the community as relationships between three entities: Context, Activities and Members. A context has a number of activities and a member can participate in a number of activities within a context. Let's consider, for example, the discussion forum. While it constitutes one context of interaction, it encompasses several activities: writing a post, reading a post, replying to a post, rating a post as shown in Figure 3 (which shows only some of the possible mapping between context and activities).

Once the social behaviour is captured, we compute the popularity, engagement and social trust values for each member as described earlier. These computed trust values along with the social behaviour of the members are used to develop the recommender system.

Our recommender system exploits the three key features of the *STrust* model with two different intents: (i) reinforcement of user behaviour and (ii) motivation for the change of user behaviour. The former is about recommending content matching an individual's interests or people with similar tastes and preferences, so that the density of interaction in the community increases (similar to collaborative filtering approaches). The latter is about recommending people and content that do not only increase the density of the interactions in the community, but also

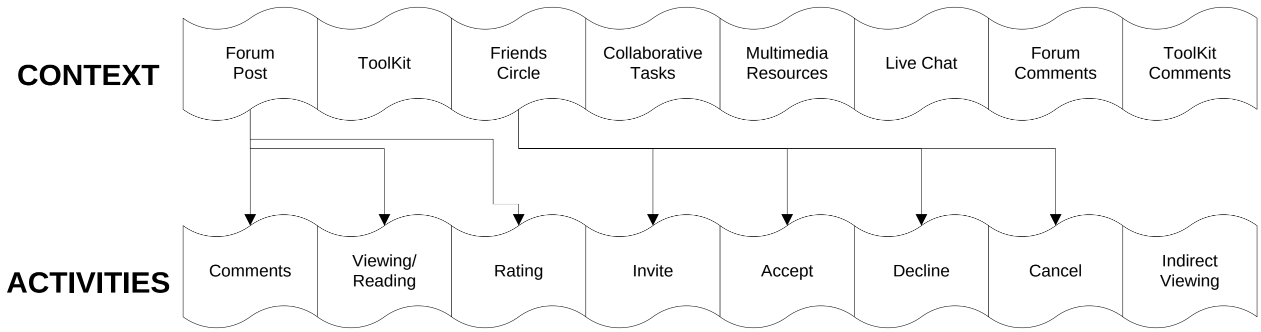

Figure 3: Context and Activities

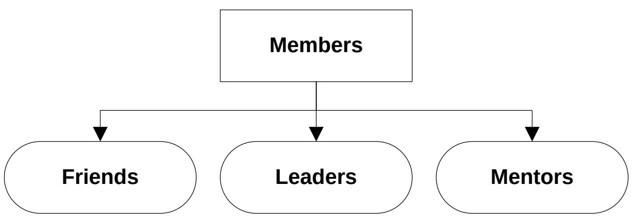

Figure 4: Types of People Recommendation in SRec

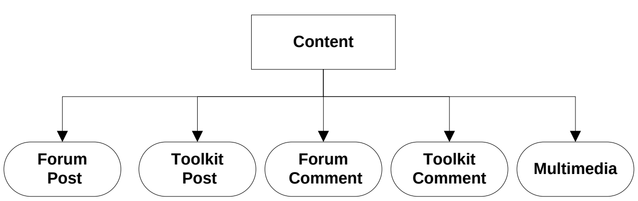

Figure 5: Types of Content Recommendation in SRec

lead users to change their behaviour (i.e., users start to like or prefer something different than what their past behaviour reflects). Thus, *SRec* recommends both people and content.

Importantly, using the different types of trust defined earlier, people can be recommended in three ways, as illustrated in Figure 4: as friends, leaders and mentors, using the different trust values (social, popularity and engagement respectively). We note that, in previous work, we validated this distinction (in particular mentors *vs* leaders) in real data sets and showed that these different roles exist in social networks, and that they are filled by different individuals [16].

Similarly, five different types of content can be recommended to community members, as illustrated in Figure 5. We next provide an overview of the mechanisms for selecting members and content for recommendation. (We refer the interested reader to [16] for the exact formulae used by the recommender system.)

Friends: A relationship exists between two members in the community in a certain context (or in the community as a whole). We exploit the member-to-member relationships to recommend two members to be friends in a certain context, using a majority rule. We recommend Bob to be a friend of Charlie if the following conditions hold:

1. The number of positive interactions of Bob is greater than a certain threshold (essentially meaning that Bob has some amount of engagement in the community).
2. The engagement trust between Bob and Charlie is greater than the overall engagement trust of Bob.
3. The number of positive interactions of Bob with Charlie is greater than the average number of positive interactions Bob has with other members.

The latter two conditions indicate that there is a priviledged relationship between Bob and Charlie. The threshold value varies from context to context and is used to control the number of recommendations. We plan to set the threshold values for different contexts through simulation and evaluate them in our community. With regards to implementation and presentation, friends are recommended in the community within the context of “Friends Circle” (as shown in Figure 3), so that members can use the same mechanisms they already used at the onset of their life in the community life to invite buddies (based on profiles, not behaviour) and send invitations to the recommended members.

Leaders: The online community provider may want to identify leaders in the community for different contexts/aspects of life. A leader is a member who has a high number of followers in the community. This typically indicates that leaders are members who provide useful posts/opinions/materials on a particular topic/context that is read/liked by many members. We capture this concept in our model using the popularity component of the social trust. A leader has a high popularity trust, although he or she might not have an overall high social trust. A leader must have a certain level of engagement (to generate followers), but a high level of engagement is not necessary.

In our online community, when leaders emerge (as computed through our recommender), we intend to have them presented (recommended) to the community through the Live Chat mechanism: leaders will be asked to conduct a Live Chat session as someone with an experience on a certain topic within a context (mainly, based on the interactions in the forum, i.e., rating, comment, etc.).

Mentors: An online community provider may want to identify a likely mentor for a community member on a certain topic. Recommendation of mentors fosters positive participation in the community and increases the social capital of the community. It is essential to have established trust relationships between a mentor and the member. In addition, it is important to determine that a mentor is a member who likes to actively engage in the community. This means the mentor must have a high level of engagement trust in the community.

When our recommender system identifies mentors in the community, we intend them to be recommended in the same way as friends are, that is within the context of “Friends Circle”. They will, however, be clearly identified as a person who could help the member to overcome their problems. Context plays an important role in selecting mentors as in leaders.

Content: Another important feature of the recommendation system is the context-based recommendation of content. This means the recommendation system should be able to alert individual community members of the presence of activities or posts that are relevant to them. Both active and passive engagement can be boosted through recommendations using the social trust. For example, the system might recommend someone to read some specific content based on their passive engagement history. Our recommender system recommends articles/posts/comments by Charlie to Bob if the following conditions hold:

1. The number of positive interactions of Bob is greater than a certain threshold.
2. The social trust between Bob and Charlie with respect to a certain context (content within the context) is greater than the overall social trust of Bob.
3. The number of positive interactions of Bob with Charlie is greater than the average number of positive interactions Bob has with other members for similar content.

SRec recommends posts or activities in a certain context x if the social trust of the user in a particular context is greater than his or her social trust in the community (over all possible contexts).

Besides recommending content to community members, we intend to use SRec to indicate to the community provider the types of informational resources that are of interest to this community. This is to help the community provider to include relevant resources and seed discussions on relevant topics so that the members remain engaged in the community. This in turn increases the social capital of the community. We will use the social trust values to provide such recommendations. For example, the community provider might be recommended to provide activities related to the context if the social trust of the community in that particular context is higher than the overall social trust in the community.

DISCUSSIONS AND CONCLUSIONS

We have presented SRec, a recommendation system based on social behavior. SRec is built upon a social trust model. SRec is able to recommend both people and content, and it distinguishes between three different types of people: friends, leaders and mentors. We believe this distinction recognises the fact that people play different roles in a community and to each other (as we validated using real data sets), and that it will add to the richness of the social recommendations that can be made in a community. Our specific community has been launched recently. We will soon be able to use our social trust model and its accompanying recommender within it, thus testing their effectiveness.

ACKNOWLEDGMENTS

We thank Payam Aghaei Pour, Nathalie Colineau, Hon Hwang, Brian Jin, Alex Sun and Bo Yan for their contributions to the community design and implementation.

This research has been partially funded under the CSIRO-Centrelink Human Services Delivery Research Alliance (HSDRA).

REFERENCES

1. Schafer, B., Konstan, J and Riedi, J.. Recommender systems in e-commerce. In *Proceedings of the 1st ACM conference on Electronic commerce* (EC '99). ACM, New York, NY, USA, (1999), 158-166.
2. Akehurst, J., Koprinska, I., Yacef, K., Pizzato, L.A.S, Kay, J. and Rej, T. Explicit and Implicit User Preferences in Online Dating. *International Workshop on Behavior Informatics,* Lecture Notes in Computer Science 7104, Springer, (2012), 15-27
3. Lops, p., Gemmis, M., Semeraro, G., Narducci, F. and Musto, C. Leveraging the linkedin social network data for extracting content-based user profiles. In *Proceedings of the fifth ACM conference on Recommender systems* (RecSys '11). ACM, New York, NY, USA, (2011), 293-296.
4. Guy, I., Ur, S., Ronen, I., Perer, A. and Jacovi, M. Do you want to know?: recommending strangers in the enterprise. In *Proceedings of the ACM 2011 conference on Computer supported cooperative work* (CSCW '11). ACM, New York, NY, USA, (2011), 285-294.

5. Balabanovi, M. and Shoham, Y. Fab: content-based, collaborative recommendation. *Communication of the ACM* 40, 3 (1997), 66-72.

6. Pazzani, M.J., Billsus, D. Content-Based Recommendation Systems. *The Adaptive Web The Adaptive Web, Methods and Strategies of Web Personalization.* Lecture Notes in Computer Science 4321, Springer (2007), 325-341.

7. Sarwar, B., Karypis, G., Konstan, J. and Riedl, J. Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th international conference on World Wide Web* (WWW '01). ACM, New York, NY, USA, (2001) 285-295.

8. Ma, H., Zhou, D. and Liu, C. Lyu, M.R. and King, I. Recommender systems with social regularization. In *Proceedings of the fourth ACM international conference on Web search and data mining* (WSDM '11). ACM, New York, NY, USA, (2011) 287-296.

9. Golbeck, J. Generating predictive movie recommendations from trust in social networks. In *Proceedings of the 4th international conference on Trust Management* (iTrust'06). Springer-Verlag, Berlin, Heidelberg, (2006), 93-104.

10. Dubbill, E.XML Watch: Finding friends with XML and RDF . www.ibm.com/developerworks/xml/library/x-foaf/index.html

11. Gilbert, E. and Karahalios, K. Predicting tie strength with social media. In *Proceedings of the 27th international conference on Human factors in computing systems* (CHI '09). ACM, New York, NY, USA, (2009), 211-220.

12. Nepal, S., Sherchan, W. and Paris, C. Building trust communities using social trust. In *Advances in User Modeling*: *UMAP 2011 Workshops*, Liliana Ardissono and Tsvi Kuflik (Eds.). Springer-Verlag, Berlin, Heidelberg (2011) 243-255.

13. Colineau, N., Paris, C. and Dennett, A. Exploring the Use of an Online Community in Welfare Transition Programs. In the *Proceedings of the 25th BCS Conference on Human Computer Interaction (HCI 2011)*—Extended Abstract, July 4-8, 2011, Newcastle Upon Tyne, UK.

14. Colineau, N., Paris, C. and Dennett, A. Capitalising on the Potential of Online Communities to Help Welfare Recipients. In the *Government and Citizen Engagement Workshop, in conjunction with the 5th International Conference on Communities and Technologies* (C&T 2011), 29 June – 2 July, 2011. Brisbane, Australia. Proceedings published in the International Reports on Socio-Informatics (IRSI), vol. 8(2), pp. 58-65.

15. Bista, S., Nepal, S. and Paris, C. Engagement and Cooperation in Social Networks: Do Benefits and Rewards Help? *11th IEEE International Conference on Trust, Security and Privacy in Computing and Communications* (IEEE TrustCom-12), UK, (2012) (to appear).

16. Nepal, S., Paris, C., Bista, S. and Sherchan, W, A Trust Model Based Analysis of Social Networks. International Journal of Trust Management in Computing and Communications, Inderscience Publishers (In press).