Selecting Items of Relevance in Social Network Feeds*

Shlomo Berkovsky, Jill Freyne, Stephen Kimani, and Gregory Smith

Tasmanian ICT Centre, CSIRO,
Hobart, Australia
firstname.lastname@csiro.au

Abstract. The success of online social networking systems has revolutionised online sharing and communication, however it has also contributed significantly to the infamous information overload problem. Social Networking systems aggregate network activities into chronologically ordered lists, Network Feeds, as a way of summarising network activity for its users. Unfortunately, these feeds do not take into account the interests of the user viewing them or the relevance of each feed item to the viewer. Consequently individuals often miss out on important updates. This work aims to reduce the burden on users of identifying relevant feed items by exploiting observed user interactions with content and people on the network and facilitates the personalization of network feeds in a manner which promotes relevant activities. We present the results of a large scale live evaluation which shows that personalized feeds are more successful at attracting user attention than non-personalized feeds.

1 Introduction

The quantity and variety of information available online has far exceeded expectations, and yet there seems to be no end to the growth and diversity of emerging online content. A recent contributor to this relentless growth is the social web which has firmly established itself as *the* platform for sharing and consuming user contributed content. While a key focus of social media is the facilitation of communication between friends, social networking systems such as Facebook and Twitter are fast becoming locations where highly valuable content is found. The volume of content produced is overwhelming and the challenge for users is to keep up with a ferociously fast changing environment and locate items of interest.

Not all relationships on social networks (SN) are equivalent and not all shared content is interesting to all users. Consider the dimensions of online social relationships; some are family based, some involve colleagues or professional connections, some reflect real world friendships and others exist exclusively online. It is natural that the strength and nature of an online relationship will influence the interest that one user has in the activities of another. In a similar vein, some users will have a preference for particular content types over others. Facebook's attempts to keep users informed of the activities of others is to summarise *all* of the performed activities of *all* of an

* This research is jointly funded by the Australian Government through the Intelligent Island Program and CSIRO. The Intelligent Island Program is administered by the Tasmanian Department of Economic Development, Tourism and the Arts. The authors thank Nilufar Baghaei, Emily Brindal, and Mac Coombe for their contribution to this work.

Joseph A. Konstan et al. (Eds.): UMAP 2011, LNCS 6787, pp. 329–334, 2011.

individuals friends in a chronologically ordered *Network Feed*. This could be an effective tool if it were not for the average users' 130 friends and their 90 pieces of contributed content per month, but in the current circumstances this list is ineffective communication medium.

This work reports on an approach to personalizing the items in a user's Network Feed in order to promote relevant updates. The personalization technique presented exploits learned user-to-user tie strength and user-to-action activity strength indicators in order to judge the relevance of each item in a network feed [2]. We report on an evaluation of our model as part of a large-scale live user study of an experimental eHealth portal. We logged all user interaction with the portal and analysed the uptake of the feeds. Initial results show that the uptake of the personalized feeds is higher than of the non-personalized ones and call for future work on the impact of feed ranking and accuracy of the relevance scoring mechanism.

2 Related Work

The sheer growth of SN contributes highly to the information overload problem, which can be only partially addressed by simple activity feeds. Hence, we have seen a move toward the development of predictive models which examine the relationships between individuals and other users, as well as content types on a social network. For the most part work in this area has concentrated on determining the predictive models rather than exploiting these models to benefit users.

Gilbert and Karaholios developed a *tie strength* model [3], which classified the strength of the relationship between users as weak or strong based on 74 Facebook factors, divided into seven categories: intensity, intimacy, duration, reciprocal services, structure, emotion, and social distance. Paek *et al.* used SVM-based classifiers to elicit a set of most predictive features and then used these features to compute the importance of activities included in Facebook news feeds [5]. The predictive models were accurate in both cases, but the factors included were specific to the type of social networking systems on which they were generated. The evaluations were conducted with small cohorts of users whereas our work reports on a large-scale evaluation.

Wu *et al.* developed a model for computing professional, personal, and overall closeness of users of an enterprise SN [6]. 53 observable SN factors were derived and divided into five categories: user factors, subject user factors, direct interaction factors, mutual connection factors, and enterprise factors. Freyne *et al.* developed a system for recommending SN activities of an interest based on long- and short-term models of content viewed and activities performed by users [2], they simulated feed personalization using offline logs, whereas our work reports on a live user evaluation.

3 Activity Relevance Score Computation

Network activity feeds present a target user T with a list of activities performed by other users of the SN. Each feed item, I, consists of at least two components: the user u_x who performed the activity and the action a_z that was performed. Typically, both the user and action are hyperlinked, facilitating access to the profile of the user who performed the activity and the content viewed/contributed by the activity (see Figure 1). The overall relevance score of the feed activity $S(T,I)$ is computed as a weighted combination of the relevance scores of the two components:

$$S(T,I)=w_1 S_U(T,u_x)+w_2 S_A(T,a_z)$$

where w_1 and w_2 denote the relative weights of the components. In our case we assign $w_1=0.8$ and $w_2=0.2$, to emphasise activities performed by relevant users. In the rest of this section we elaborate on the computation of user-to-user $S_U(T,u_x)$ and user-to-action $S_A(T,a_z)$ relevance scores.

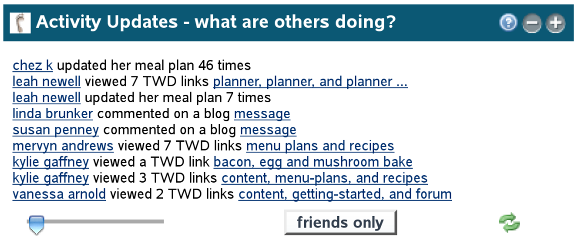

Fig. 1. Example activity feed

For the computation of the user-to-user friendship score $S_U(T,u_x)$, we adopt the weighting model developed by Wu *et al.* in [6]. The enterprise factors are inapplicable to our eHealth portal and, therefore, we use four categories of factors:

- User factors (UF) – online behaviour and activity of the target user T.
- Subject user factors (SUF) – online behaviour and activity of the subject user u_x.
- Direct interaction factors (DIF) – direct communication between T and u_x.
- Mutual connection factors (MCF) – communication between T and $\{u_y\}$ and between u_x and $\{u_y\}$, where $\{u_y\}$ is the set of common friends of T and u_x.

The user-to-user relevance score $S_U(T,u_x)$ is computed as a weighted combination of the scores of the four categories:

$$S_U(T,u_x)=w_3 S_{UF}(T,u_x)+w_4 S_{SUF}(T,u_x)+w_5 S_{DIF}(T,u_x)+w_6 S_{MCF}(T,u_x)$$

Since the features of the system presented in [6] were similar to those provided by our eHealth portal, we assign to the four categories relative weights proportional to the weights derived there: $w_3=0.178$, $w_4=0.079$, $w_5=0.610$, and $w_6=0.133$.

Category scores $S_{UF}(T,u_x)$, $S_{SUF}(T,u_x)$, $S_{DIF}(T,u_x)$, and $S_{MCF}(T,u_x)$ are computed as a weighted combination of the scores of the factors in each category. Overall, we derived 32 factors for the UF and SUF categories and 28 factors for the DIF and MCF categories. The score for each factor is computed based on the observed user interaction with the SN and normalised to the $[0,1]$ range. Table 1 presents the four factors with the highest relative weights within each category. Note that SUF factors are identical to UF, but the score is computed for u_x rather than for T. Similarly, MCF factors are identical to DIF, but the score is averaged across all mutual friends $\{u_y\}$.

The frequency of performing actions is considered the main indicator of user-to-action relevance. We denote by $f(T,a_z)$ the frequency of user T performing action a_z, by $f(T)$ the average frequency of all actions performed by T, by $f(a_z)$ the average frequency of all users performing a_z, and by $f()$ the average frequency of all actions performed by all users. The user-to-action relevance score $S_A(T,a_z)$ is computed by

$$S_A(T, a_z) = \frac{f(T, a_z)}{f(T)} / \frac{f(a_z)}{f()}$$

Hence, we first computed the relative relevance of a_z for T and then normalised it by the relevance of a_z for all users.

Table 1. User-to-user relevance factors and their weights

UF		SUF		DIF		MCF	
factor	weight	factor	weight	factor	weight	factor	weight
# forum posts added by *T*	*0.02031*	# forum posts added by u_x	*0.00899*	Has *T* friended u_x	*0.07627*	Has *T* friended $\{u_y\}$	*0.01656*
# posts in *T*'s blog	*0.02031*	# posts in u_x's blog	*0.00899*	# days *T* interacted with u_x	*0.04576*	Average # days *T* interacted with $\{u_y\}$	*0.00994*
# *T's* comments in blogs of others	*0.01015*	# u_x*'s* comments in blogs of others	*0.00449*	# *T*'s posts in u_x's blog	*0.03814*	Average # posts in members of $\{u_y\}$'s blog	*0.00828*
# images in *T*'s profile	*0.01015*	# images in u_x's profile	*0.00449*	# mutual friends of *T* and u_x	*0.02670*	Average # mutual friends of *T* and members of $\{u_y\}$	*0.00580*

4 Evaluation

We evaluated the feed relevance prediction computation in a live study involving users of an experimental eHealth portal. 2,813 users participated in the study for a period of 12 weeks, from September to November of 2010. The users were randomly divided into several groups, such that about half of them were exposed to personalized and half to non-personalized activity feeds. No personalized feeds were provided during the first week, due to the bootstrapping phase required to determine the relevance scores. From week 2 onwards, users in the personalized groups were exposed to personalized feeds, as presented in Section 3. Users in the non-personalized groups were presented with activities ordered in reverse chronological order. User-to-user and user-to-action scores were calculated offline nightly. Figure 1 depicts the feed interface. By default, the feeds presented 20 activities (with the highest relevance score or most recent timestamp), but users could adjust the number of items.

The results below address our initial findings that focus on the uptake of the feeds, as reflected by the observed click-through rate. Table 2 summarises the number of users, sessions, sessions with feed clicks, feed clicks, and two computed click-through rates (CTR_u – number of clicks per user and CTR_s – number of clicks per session with feed clicks), as observed for both groups from week 2 onwards. As can be seen, the uptake of the personalized feeds was higher than that of the non-personalized feeds.

This is supported by observing the CTR_s score over time. Figure 2 depicts the average CTR_s in both groups for the entire duration of the study. Only days, for which four clicks or more were logged for each group, are included. Due to the variability of the CTR_s values (spikes of activity were originated by email reminders sent to all users and by bursts of SN activity that affected both groups of users similarly), we plotted also the logarithmic regression curves. The uptake curve of the personalized feeds was steadily superior to the curve of the non-personalized feeds. The difference between the groups was statistically significant, $p=0.0195$. This indicates that the activities in the personalized feeds were deemed more relevant than in the non-personalized ones.

Table 2. Overall uptake of the feeds

	users	sessions	$sessions_{cl}$	clicks	CTR_u	CTR_s
personalized	*1,397*	*12,193*	*390*	*901*	*0.6450*	*2.3103*
non-personalized	*1,416*	*11,386*	*382*	*805*	*0.5685*	*2.1073*

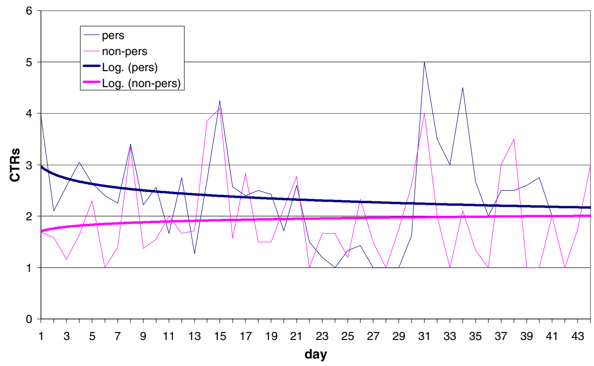

Fig. 2. Feed click-through rates over time

We also computed the number of sessions that included multiple feed clicks. Table 3 summarises the number of feed clicks in a session and the number of sessions that included this (or higher) number of clicks, as was observed for both groups. As can be seen, the number of sessions with multiple clicks in the personalized group was higher than in the non-personalized group. The difference between the groups, however, was not statistically significant. This strengthens our conclusion that the activities in the personalized feeds were deemed more relevant than in the non-personalized feeds.

Table 3. Multiple feed clicks in a session

number of clicks	2	3	4	5	6	7	8	9	10	11	12	13	14	15
personalized	*154*	*87*	*55*	*45*	*33*	*30*	*25*	*20*	*16*	*15*	*9*	*6*	*4*	*4*
non-personalized	*157*	*82*	*53*	*36*	*20*	*14*	*11*	*10*	*7*	*5*	*4*	*2*	*2*	*1*

Finally, we computed the user-to-user, user-to-action, and overall relevance scores of the clicked activities, as evolved for the entire duration of the study. Figure 3 depicts the average relevance scores of the clicked feeds items, again only for days with four or more clicks. Once again, we plotted the logarithmic regression curves due to the variability of the relevance scores.

The user-to-action relevance scores stabilise rapidly and remain stable for the entire duration of the study. However, the user-to-user relevance scores steadily increase over time and, as a result, overall relevance scores, which are dominated by the user-to-user relevance, also increase over time. This is in line with earlier personalization research findings, which indicate that the accuracy of personalization improves as the amount of information available about the users increases.

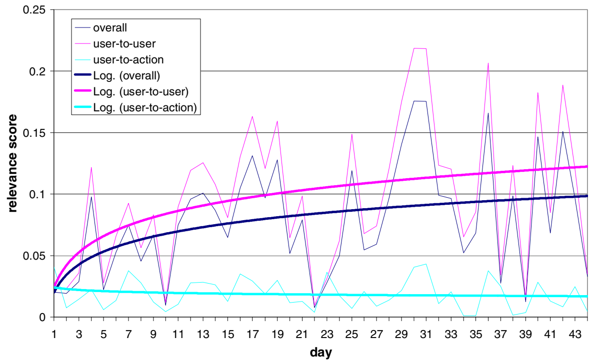

Fig. 3. User-to-user, user-to-action, and overall relevance scores over time

5 Conclusions and Future Work

This work was motivated by the growing volume of information included in SN activity feeds. We developed a personalized model for predicting the relevance score of each activity in the feed and re-ranking the feed accordingly. The model was evaluated with a large set of users for an extensive period of time. Initial analyses of the results showed that the uptake of the personalized feeds is higher than of the non-personalized ones and that the relevance of the clicked activities increases over time.

In future analyses, we plan to address the impact of ranking on the uptake of activities. High-ranked activities are normally clicked more frequently than low-ranked ones and we will revise our model to factor out the effect of ranking. Also, we will revisit the weighting model adopted from [6], and will derive a new model appropriate for the eHealth application domain in general and our portal in particular.

References

1. Facebook Statistics, `http://www.facebook.com/press/info.php?statistics` (accessed January 2011)
2. Freyne, J., Berkovsky, S., Daly, E.M., Geyer, W.: Social Networking Feeds: Recommending Items of Interest. In: Proceedings of RecSys, Barcelona (2010)
3. Gilbert, E., Karahalios, K.: Predicting Tie Strength with Social Media. In: Proceedings of CHI, Boston (2009)
4. Noakes, M., Clifton, P.: The CSIRO Total Wellbeing Diet. Penguin Publ. (2005)
5. Paek, T., Gamon, M., Counts, S., Chickering, D.M., Dhesi, A.: Predicting the Importance of Newsfeed Posts and Social Network Friends. In: Proceedings of AAAI, Atlanta (2010)
6. Wu, A., DiMicco, J.M., Millen, D.R.: Detecting Professional versus Personal Closeness using an Enterprise Social Network Site. In: Proceedings of CHI, Atlanta (2010)