Personalized Network Updates: Increasing Social Interactions and Contributions in Social Networks

Shlomo Berkovsky[1,2], Jill Freyne[1], and Gregory Smith[1]

[1] CSIRO ICT Centre,
Marsfield, NSW, Australia
firstname.lastname@csiro.au
[2] NICTA,
Alexandria, NSW, Australia
firstname.lastname@nicta.com.au

Abstract. Social networking systems originally emerged as tools for keeping up with the daily lives of friends and strangers. They have established themselves as valuable resources and means to satisfy information needs. The challenge with information seeking through social networks is that their immense success and popularity is also a weakness. The data deluge facing users has surpassed comfortably managed levels and can impact on the quality and relevance of the information consumed. We developed a personalized model for predicting the relevance of news feed items, in order to facilitate personalized feeds. Results of a live analysis show that our approach successfully identifies and promotes relevant feed items, with the knock-on effects of increasing interaction between users and the contribution of user generated content.

1 Introduction

Social media refers to online services and portals that foster user interaction and a sense of community, allowing their users to establish and maintain relationships, share information, and express opinions. Social media evolved into a class of applications that build on the foundations of Web 2.0 and allow the creation and exchange of user-generated content [9]. Among other social media applications, social networking sites (or, in short, social networks – SNs) have gained remarkable popularity, and are fast becoming locations, where content is shared and found. Facebook alone reports more than 800 million active users (more than half of which log on daily), with an average user connected to 80 communities and events, having 130 friends, and using the system for about one hour a day [4].

The volume of content generated by SN users is enormous, and there seems to be no foreseen limits to the growth and diversity of this content. The initial mechanism devised for keeping users abreast of the activities of others is the News or Activity Feed, a reverse chronologically ordered list showing the activities of friends or followees (see two popular examples in Figure 1). The feed typically communicates updates and activities carried out by all of a user's friends and followees in one list. While simple and easy to understand, the feed was not designed to cope with the huge number of friends and followees, or the volume and diversity of content contributed nowadays to SNs, and is crumbling under the pressure being placed upon it.

J. Masthoff et al. (Eds.): UMAP 2012, LNCS 7379, pp. 1–13, 2012.

Fig. 1. Example news feeds: Facebook (left) and Twitter (right)

In response to this growing issue, we investigate in this work the use of personalization algorithms, which are common solutions in other information overload situations, in order to identify the content of news feeds that is most valuable for each user. We extend earlier works [7,14] and exploit the observable activities of SN users to predict the degree of relevance of the feed items for users. We judge the relevance of the feed items using two factors: user-to-user relationship strengths and user-action interest scores. Then, we filter the feed to separating relevant news items from noise.

This paper follows up on earlier work [1] and evaluates the developed feed personalization approach as part of a large-scale live user study of an experimental eHealth portal. We present an extensive analysis of the uptake of the feeds, as well as of user interactions resulting from the feed clicks. The results show that the personalization successfully highlights relevant SN activities, assists users in establishing and maintaining online friendships, and increases contribution of content (wall comments and blog posts). Hence, the pivotal contribution of this work is the thorough investigation of the impact of personalized activity feeds on user behaviour on a SN site, in particular on content contribution and friending.

The rest of this paper is structured as follows. Section 2 surveys related work on personalization of SN feeds. Section 3 provides an overview of the Online TWD Portal. Section 4 presents our feed personalization algorithms. Section 5 presents the experimental evaluation and discusses the obtained results. Finally, Section 6 concludes the paper and outlines our future research directions.

2 Related Work

To facilitate a reorganisation of the news feed, a robust mechanism capable of judging the relevance of the users and actions in each feed item is required [3]. Recently, we have seen works on predictive models that examine the relationships between users and items on SNs, which are ideal for this purpose. Although these works concentrate on the development of the models, they pave the way to the application of these models in personalized algorithms, in order to alleviate the data deluge facing SN users.

Gilbert and Karaholios developed the tie strength model [7], which classified the strength of a relationship between Facebook users as weak or strong based on 74 factors, divided into seven categories: intensity, intimacy, duration, reciprocal service, structure, emotion, and social distance. Paek *et al.* used SVM-based classifiers to

elicit a set of most predictive features and then exploited these features to compute the importance of activities included in Facebook news feeds [12]. These works evaluated their predictive models with small cohorts of users and, although the models were accurate in both cases, the factors included in the models were specific to the SN system on which they were generated.

Wu *et al*. developed a model for computing professional, personal, and overall closeness of users of an enterprise SN [14]. 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 an approach for recommending SN activities of interest based on long- and short-term models of content viewing and activities performed by users [5]. They simulated feed personalization using offline logs and simulated the events from these logs. Guy *et al.* proposed to consider the content of the activity feeds for profiling users of an enterprise SN [8]. These works judged the strength or closeness of users based on their online behaviour, observed interactions, and content of their feeds. Our work expands the existing models to include activity interest and illustrates a high-value use case of personalized news feeds in a live SN.

3 The Online Total Wellbeing Diet Portal

The evaluation of the proposed feed personalization approach was conducted as part of a large-scale user study of an experimental eHealth portal. The portal aimed to support people embarking of the CSIRO Total Wellbeing Diet (TWD) program [11] and contained dietary information and tools, as well as typical SN functionalities (see Figure 2). The online information mirrored the content of the TWD book and included recipes, exercises, menu plans, shopping lists, and other health links. In addition, the portal included several tools, such as a meal planer and weight tracker, which provided users with real-time feedback on decisions and progress [6].

The goal of the SN was to provide online mutual support for dieters. Each user was represented by a profile page, which contained personal information, an image

Fig. 2. TWD Online portal

gallery, a personal message board (wall) and a blog. The blogs were free-text diaries, to which the users could contribute as often as they wished. Privacy restrictions of the blogs were set by their owners, but public blogs could be seen by any user. To facilitate community-based information sharing, the portal contained a discussion forum. Here, the users could ask questions, provide support, seek advice, and discuss ideas and thoughts with the community. The forum was monitored by domain experts, who answered health-, exercise-, and nutrition-related questions posted by users.

A key goal of the portal was to support and encourage dieters by exposing them to the thoughts and activities of others on the diet. By highlighting activities, like meal planning, weighing in, browsing recipes and exercises, and reading/writing blogs, we aimed to encourage users to also carry out these activities. By showing the users who carries out the activities, we introduced them to like-minded people on which they can call for support. Similarly, by providing links to blog and forum, we allowed users to provide feedback, support each others, and be inspired to contribute. To this end, it was important to make the activities and contributions of others highly visible. Hence, our portal included an activity feed, which aggregated the interactions of users with the content, tools, and SN (see the interface of the feed in Figure 3).

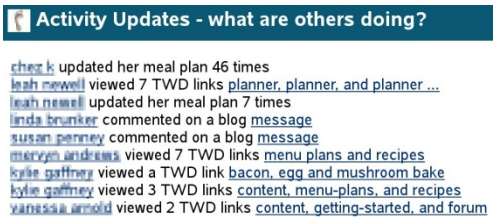

Fig. 3. SN activity feed

4 Personalized News Feeds

The feed presented a target user u_t with a list of activities performed by others. Each item i_x included in the feed references two components: the subject user u_x who performed the activity and the action a_x that was performed, e.g., wall comments, forum posting, or content viewing. When the feed was visualised, both the user name and action were hyperlinked, such that clicks on u_x, i.e., *user clicks*, provided access to the profile page of the user who performed the activity, whereas clicks on a_x, i.e., *action clicks*, led to the content viewed/contributed by the activity.

Our personalization algorithm assigns to each feed item a relevance score $S(u_t,i_x)$, which represents the predicted level of interest that a target user u_t will have in item i_x and is computed as a weighted linear combination of a user-to-user score $S_u(u_t,u_x)$ and a user-action score $S_a(u_t,a_x)$, where w_u and w_a denote the relative weights of the two components as seen in Equation 1:

$$S(u_t,i_x)=w_uS_u(u_t,u_x)+w_aS_a(u_t,a_x) \quad (1)$$

Following the feed mechanisms applied by the popular SNs (see examples in Figure 1), we presumed that activities of users with which u_t had closer relationships

would attract higher interest than activities involving actions of importance to users. Hence, we assigned static weights of $w_u=0.8$ and $w_a=0.2$, which emphasise activities performed by relevant users.

The user-to-user relevance score $S_u(u_t,u_x)$ reflects the closeness of a target user u_t and a subject user u_x, derived solely from their online interactions. To compute this relevance score, we deployed a modified variant of the tie strength model developed in [7], adapted according to the closeness factors proposed in [14]. Some of the original factors of [14] were related to the enterprise environment and were found inapplicable to the TWD Online portal. Hence, we used four categories of factors:

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

The user-to-user relevance score $S_u(u_t,u_x)$ was computed as a linear combination of the scores of these four categories of factors:

$$S_u(u_t,u_x)=w_{uf}S_{uf}(u_t,u_x)+w_{suf}S_{suf}(u_t,u_x)+w_{dif}S_{dif}(u_t,u_x)+w_{mcf}S_{mcf}(u_t,u_x) \quad (2)$$

Since the functionality and the components of the enterprise SN presented in [14] were similar to those offered by the TWD Online portal, we assigned to the scores of these four categories relative weights that are proportional to the original weights derived in [14]: $w_{uf}=0.178$, $w_{suf}=0.079$, $w_{dif}=0.610$, and $w_{mcf}=0.133$.

Category scores $S_{uf}(u_t,u_x)$, $S_{suf}(u_t,u_x)$, $S_{dif}(u_t,u_x)$, and $S_{mcf}(u_t,u_x)$ were computed as a linear combination of the scores of the factors belonging to each category. For the UF and SUF categories, we derived 32 factors that reflect the individual behaviour of u_t and u_x,. These include the number of forum/blog/wall posts they initiated/answered/rated, the number of active sessions/days, the number of times they updated/viewed the content/images in their profiles, and others. Also, we derived 28 factors for the DIF and MCF categories that, respectively, reflect the direct interaction between u_t and u_x, and their interaction with the set of their mutual friends $\{u_y\}$, i.e., the users who friended both u_t and u_x. These factors included the number of answers/ratings to each other's forum/blog/wall posts, number of sessions/days they interacted with each other, whether they friended each other, duration of their friendship, and others.

The scores of the factors were computed using the observed frequencies of various user interactions with the TWD Online portal and normalised[1] to the $[0,1]$ range. The scores of the UF and SUF factors were computed in the same manner, but using the frequencies observed for users u_t and u_x, respectively. The scores of the DIF factors were computed using the frequencies of direct interactions between the two users. The scores of the MCF factors were computed[2] by averaging the individual DIF scores computed for the two users u_t and u_x across their mutual friends, i.e., across a set of users $\{u_y\}$, who are online friends of both u_t and u_x. Table 1 presents four factors with the highest weight within each category.

[1] The scores of the UF and SUF factors were normalised by dividing the observed frequency of the user by the maximal frequency observed for any other user. The scores of the DIF and MCF factors, which involve multiple users, were normalised using Jaccard's similarity coefficient. Due to space limitations, the details of the normalisation are omitted.

[2] Due to space limitations, the details of this computation are also omitted.

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 u_t	*0.02031*	# forum posts added by u_x	*0.00899*	has u_t friended u_x	*0.07627*	has u_t friended $\{u_y\}$	*0.01656*
# posts in u_t's blog	*0.02031*	# posts in u_x's blog	*0.00899*	# days u_t interacted with u_x	*0.04576*	# days u_t interacted with $\{u_y\}$	*0.00994*
# u_t*'s* comments in blogs of others	*0.01015*	# u_x*'s* comments in blogs of others	*0.00449*	# u_t's posts in u_x's blog	*0.03814*	# posts in $\{u_y\}$'s blog	*0.00828*
# images in u_t's profile	*0.01015*	# images in u_x's profile	*0.00449*	# mutual friends of u_t and u_x	*0.02670*	# mutual friends of u_t and $\{u_y\}$	*0.00580*

In a similar manner, we calculated the user-action interest score $S_a(u_t,a_x)$. This score reflects the importance of action a_x for user u_t and is informed by the frequency of performing the action a_x and the frequencies of performing other actions [2]. The user-action relevance score $S_a(u_t,a_x)$ is calculated as shown in Equation 3, where $f(u_t,a_x)$ is the frequency of user u_t performing a_x, $f(u_t)$ is the average frequency of all actions performed by u_t, $f(a_x)$ is the average frequency of all users performing a_x, and $f()$ is the average frequency of all actions performed by all users.

$$S_a(u_t, a_x) = \frac{f(u_t, a_x)}{f(u_t)} / \frac{f(a_x)}{f()} \tag{3}$$

This computation quantifies the relative importance of a_x for u_t and normalises it by the relative importance of a_x for all users. The user-action score $S_a(u_t,a_x)$ computed using Equation (3) and the user-to-user score $S_u(u_t,u_x)$ computed using Equation (2) are aggregated into the overall feed item score $S(u_t,i_x)$, as shown by Equation (1). Items having the highest predicted scores were included in the feed.

5 Evaluation

Over 8000 individuals were recruited to participate in a large scale study of the TWD Online portal over a period of 12 weeks in late 2010. The study mainly focused on health-related outcomes, such as weight loss and engagement with the diet. 5279 users participated in the study, but only a portion of these were relevant to the presented analyses, as not all users had access to the activity feeds and not all those who had access interacted with the portal when the personalization became active (after a bootstrapping period of one week). Of those who had access to the feeds, each user was randomly allocated to an experimental group at recruitment time, such that half were exposed to personalized and half to non-personalized feeds. Users in the personalized group were shown personalized feeds, in which the items were scored as described in Section 3, while users in the non-personalized control group were presented

with chronologically ordered feeds. By default, the feeds included 20 items (with the highest scores or most recent timestamps), but the users could adjust this parameter.

5.1 Activity Feed Uptake

Overall, the level of user interaction with the news feed was lower than expected, with 137 users generating 530 feed clicks over the course of the study[3]. Table 2 summarises the number of users who interacted with the feed, i.e., clicked on feed items, the number of sessions that included feed clicks, the overall number of logged clicks and their breakdown into user clicks (on the user name) and action clicks (on the action that was carried out), and the user- and session-based click through rates, CTR_u and CTR_s [13], computed as the ratio between the overall number of clicks and the number of users and sessions, respectively, as observed for both groups.

Table 2. Feed uptake

	users	sessions	$clicks_u$	$clicks_a$	clicks	CTR_u	CTR_s
personalized	*64*	*125*	*159*	*87*	*246*	*3.844*	*1.968*
non-personalized	*73*	*159*	*155*	*129*	*284*	*3.890*	*1.786*

We note that users interacted more with the non-personalized feeds (both the number of users and overall numbers of clicks), but the observed CTR_u was comparable. However, the personalized feeds appear to provide more relevant information, as communicated by their higher session-based CTR_s. Note that the percentage of user clicks in the personalized feeds was 64.6%, in comparison 54.6% in the non-personalized feeds. That is, users in the personalized group were more interested in the subject users who performed the activities than users in the non-personalized group. This can be explained by the weighting mechanism of Equation (1), which assigned 80% of the overall weight to $S_u(u_t,u_x)$ and only 20% to $S_a(u_t,a_x)$. Thus, activities of users with high user-to-user score dominated over activities with high user-action score, and this was reflected by the higher percentage of user clicks.

We compared the average number of feed clicks per session with clicks for each group. Table 3 shows the percentage of sessions, in which N (from 2 to 12) feed clicks or more occurred. For example, in 42.4% of sessions where personalized feeds were presented and at least one feed click was logged, two or more feeds clicks were recorded, while for the non-personalized feeds two or more clicks were recorded in 37.1% of sessions. The percentage of sessions with multiple clicks in the personalized group was consistently higher than those observed in the non-personalized group, although no statistically significant difference was detected[4]. Overall, the higher CTR_s and rate of sessions with multiple clicks show that the personalized feeds attracted more user attention and, thus, deemed more relevant than the non-personalized feeds.

[3] A small portion of clicks were omitted from the analyses due to technical issues that resulted in unreliable user logs.

[4] All statistical significance results refer to a two tailed t-test assuming equal distribution.

Table 3. Multiple feed clicks in a session

N – number of clicks	2	3	4	5	6	7	8	9	10	11	12
personalized [%]	*42.4*	*20.0*	*15.2*	*12.8*	*8.0*	*7.2*	*5.6*	*5.6*	*4.0*	*4.0*	*1.6*
non-personalized [%]	*37.1*	*17.6*	*9.4*	*7.6*	*6.3*	*4.4*	*3.1*	*1.8*	*1.2*	*1.2*	*1.2*

5.2 Feed Ranking

The motivation for personalizing activity feeds was to assist users in finding relevant information by re-ranking the feed and promoting items in which they are likely to be interested. Hence, we compared the distribution of the clicked items across the two feeds, to verify whether the personalization affected the positioning of user clicks within the feeds [13]. The average rank of the clicked items in the personalized feeds was 4.00 in comparison to 4.46 (smaller numbers indicate higher rank) in the non-personalized feeds, and the difference was statistically significant, $p<.05$, illustrating the influence of personalization in promoting relevant feed items. We computed the average rank of the clicked items, as observed for weeks 2 to 9 (week 1 was bootstrapping and weeks 10-12 did not attract sufficient clicks). As can be seen in Figure 4, the personalized feeds consistently outperformed the non-personalized ones and presented relevant items closer to the top of the feed. The difference between the two groups was statistically significant, $p<.05$.

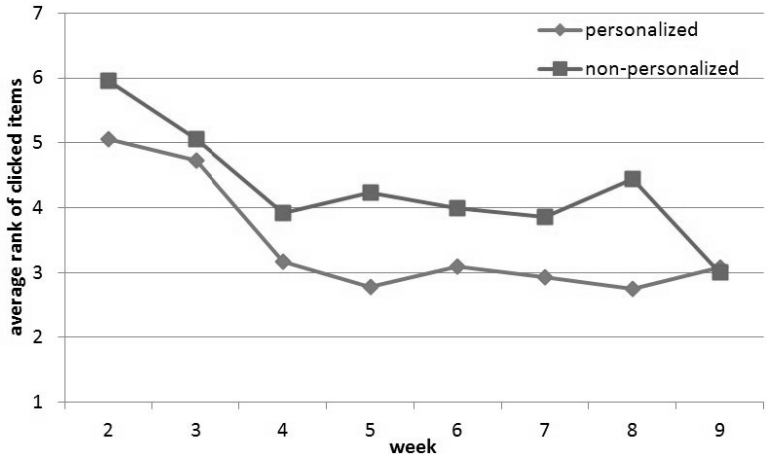

Fig. 4. Average rank of the clicked items over time

We also assessed the impact of personalization on the distribution of feed clicks. Figure 5 details the percentage of clicks observed for positions 1 to 10 in the feed for both groups. In the non-personalized feeds, a comparable distribution of 10-15% across the top seven positions was observed. Thus, users showed a similar level of interest in a large group of items, rather than only in the top-ranked most recent items. In the personalized feeds, we note a preference for items in the top position,

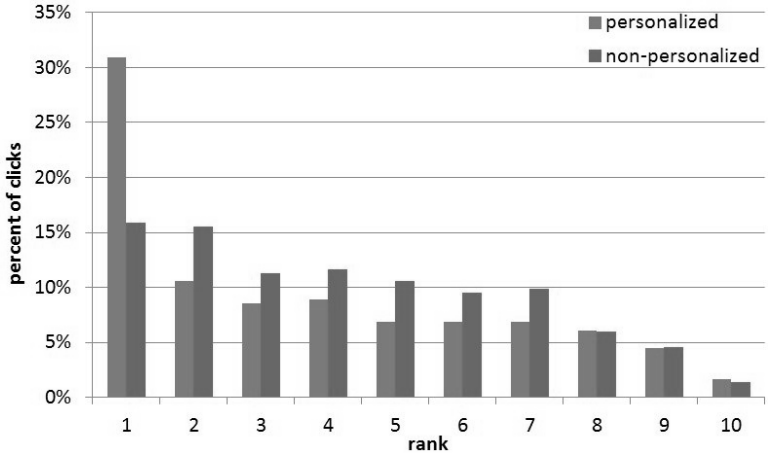

Fig. 5. Rank of the clicked items

which attracted 30.9% of clicks – double that of the top position in the non-personalized feeds. The observed click patterns clearly show that users found the personalized top items to be more relevant than the non-personalized ones. Hence, the personalization successfully promoted relevant items and making relevant information easier to find.

5.3 Impact on User Activities

Another aim of the activity feed was to highlight activities of others, in order to spark social relationships and social learning. To understand the extent to which this was achieved, we examined the activities carried out by users in the sessions where they interacted with the activity feed. Focusing on the social interactions sparked by the feed, we examined the number of blog and forum activities, profile views, wall postings, and the overall session length observed in each session that included feed clicks. In addition, we split the blog and forum activities into contribution activities (posting, responding) and consumption activities (viewing), as shown in Table 4.

Table 4. Sessions with feed clicks

	blog	*bl-cont*	*bl-cons*	forum	*f-cont*	*f-cons*	profile	wall	session
personalized	*10.35*	*5.54*	*4.81*	*5.58*	*1.00*	*4.58*	*7.35*	*0.71*	*48.19*
non-personalized	*9.77*	*3.92*	*5.84*	*5.04*	*0.28*	*4.76*	*4.16*	*0.31*	*45.12*

The personalization increased (although not significantly) the overall session length from 45.12 to 48.19 activities. This was reflected by more blog, forum, profile viewing, and wall posting activities. Out of all these, the profile viewing and wall posting were statistically significant, both $p<.01$. It must be noted that these two activities were not independent, as the wall is located on a user's profile page, such that all wall postings are preceded by a profile view. That said, less than 10% of profile

views resulted in wall posts, showing that users view profiles not only when they intend to write wall messages. Furthermore, significant increases in the contribution to blogs and forums were made by users viewing personalized feeds, when compared to non-personalized feeds. Thus, personalization sparked social interactions with other users and encouraged contribution of content to blogs and forums, as intended.

Refining this analysis and concentrating on activities immediately following the feed clicks, Table 5 summarizes the average number of SN-related activities recorded within five user activities following a feed click[5]. We note similar trends: higher levels of blog and forum contributions, profile views, and wall posts within the immediate activities following clicks. The most frequent activity to follow a click was profile viewing. Note that in this case, the majority of forum and blog activities were consumption rather than contribution. This is in line with prior research, which showed that consumption of user-generated content normally exceeds contribution [10].

Table 5. Five activities following feed clicks

	blog	*bl-cont*	*bl-cons*	forum	*f-cont*	*f-cons*	profile	wall
personalized	*0.760*	*0.187*	*0.573*	*0.407*	*0.041*	*0.366*	*1.825*	*0.037*
non-personalized	*0.630*	*0.095*	*0.535*	*0.394*	*0.007*	*0.387*	*1.429*	*0.018*

Overall, the personalization of activity feeds had a prolific impact on activities carried out on the TWD Online portal. It increased the volume of traffic to the user profiles, social interactions through message walls, and contribution of user-generated content to blogs and forum. Thus, personalization played an important role in the sustainment of the social features of the portal. The contribution of user-generated content is pivotal for the sustainability of any SN, as it invites users to return for further interactions with the portal and increases user engagement. Likewise, it is important in facilitating social support for users embarking on the diet, which, from the health perspective, was the primary goal of the TWD Online portal.

5.4 Feeds and Friending

The developed relevance scoring mechanism inherently presumes that users are more interested in the information pertaining to the actions of their articulated group of online friends in preference to information relating to other portal users. To this end, we examined the links between online relationships and feed clicks. It should be noted that the TWD Online portal is not a typical friendship-based SN reflecting offline friendships; very few users were familiar with other participants prior to the study.

Of the 246 clicks logged in the personalized feeds, in 79 cases (or 32.1%) the target and subject user established an online friendship over the course of the study. In the non-personalized feeds, this happened in 78 cases out of the 284 logged clicks (or 27.4%). Thus, the ratio of clicks on items representing activities of friends in the personalized feeds was higher than in the non-personalized feeds. We hypothesise that

[5] Feed clicks themselves were excluded and only five activities after the click were analysed. The sum of each row is less than 5, as health- and content-related activities were excluded.

the effect of the personalized ordering of feeds brought the actions of friends to the attention of users by promoting them to the top of the activity feeds.

This finding brings to the fore the issue pertaining to the 'personalization bubble'. In other words, does the personalization and its inherent focus on promoting the actions of those a user is close to limit the user's awareness of activities of other SN users? To answer this question, we measured the ratio of feed items involving activities of friends. Note that this ratio refers to the entire feed presented to a user rather than to clicked items only, which were addressed in previous analysis. Overall, the ratio of friend activities in the feeds increased over time in both groups. This is explained by the observed user attrition and skewed distribution of the established friendships. The attrition of users was stable over the course of the study: from week 2 onwards, about 20% of users who interacted with the portal in any given week, never returned to the portal in the following weeks. On the contrary, close to half of all the observed friendships were established in the first week only and the friending rate steadily declined afterwards. Thus, the density of the friendship network increased over time, as well as the percentage of feed items including activities performed by friends.

Figure 6 details the percentage of feed items involving activities of friends, as observed for weeks 2 to 9 for both groups. In the non-personalized feeds, the ratio between the number of friend and non-friend activities increased steadily from 3.1% in week 2 to 25% in week 9. However, in the personalized feeds this ratio did not increase and hovered around the 20% mark until week 6. Hence, the feeds still included closely to 80% of activities performed by non-friends and did not severely limit the user's awareness of activities of others. Afterwards, the friendship network became dense and the percentage of friend activities in the personalized feeds increased.

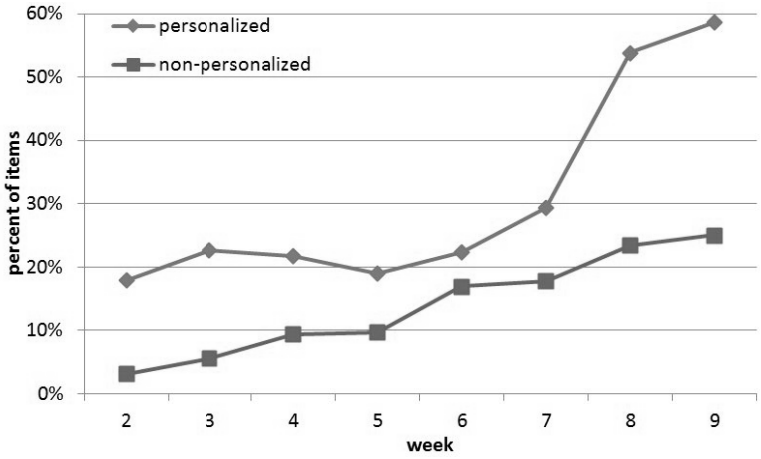

Fig. 6. Ratio of friend activities in the feeds over time

6 Conclusions and Future Work

This work was motivated by the aggravating information overload problem in SNs, which is only exacerbated by the simplistic nature of network news feeds. We developed a personalized model for predicting the relevance of news feed items to an individual and applied this model to produce personalized news feeds to participants of a live study of an online diet portal. This paper discusses the observed impact of the personalization on user interactions with the portal and other participants.

The results show that the uptake of the personalized feeds was higher than of the non-personalized ones, which had a prolific impact on the sustainability of the SN. Firstly, the personalized feeds appears to promote items of higher relevance within the news feeds, assisting the users in the identification of relevant activities. Secondly, it was found that the personalized feeds increased the contribution of user-generated content to the forum, blogs, and walls. This was observed both for immediate activities following the feed clicks and for the entire duration of the sessions. Thirdly, the personalized feeds highlighted the activities performed by online friends, while not limiting user awareness of activities of other SN users.

In the future, we plan to revise the scoring model and investigate the appropriateness of its adaptation to the domain and application in hand. We also plan to ascertain the accuracy of the user-action scoring, as it was not based on extensive prior research. We intend to extensively evaluate these in a large-scale live user study.

Acknowledgment. This research is jointly funded by the Australian Government through the Intelligent Island Program and CSIRO. The authors thank Stephen Kimani, Nilufar Baghaei, Emily Brindal, and Mac Coombe for their contributions to this work.

References

1. Berkovsky, S., Freyne, J., Kimani, S., Smith, G.: Selecting Items of Relevance in Social Network Feeds. In: Konstan, J.A., Conejo, R., Marzo, J.L., Oliver, N. (eds.) UMAP 2011. LNCS, vol. 6787, pp. 329–334. Springer, Heidelberg (2011)
2. Bohnert, F., Zukerman, I., Berkovsky, S., Baldwin, T., Sonenberg, L.: Using Collaborative Models to Adaptively Predict Visitor Locations in Museums. In: Nejdl, W., Kay, J., Pu, P., Herder, E. (eds.) AH 2008. LNCS, vol. 5149, pp. 42–51. Springer, Heidelberg (2008)
3. De Choudhury, M., Counts, S., Czerwinski, M.: Identifying Relevant Social Media Content: Leveraging Information Diversity and User Conginition. In: Proceedings of HT, Eindhoven (2011)
4. Facebook Statistics, `http://www.facebook.com/press/info.php?statistics` (accessed December 2011)
5. Freyne, J., Berkovsky, S., Daly, E.M., Geyer, W.: Social Networking Feeds: Recommending Items of Interest. In: Proceedings of RecSys, Barcelona (2010)
6. Freyne, J., Berkovsky, S., Baghaei, N., Kimani, S., Smith, G.: Personalized Techniques for Lifestyle Change. In: Peleg, M., Lavrač, N., Combi, C. (eds.) AIME 2011. LNCS, vol. 6747, pp. 139–148. Springer, Heidelberg (2011)

7. Gilbert, E., Karahalios, K.: Predicting Tie Strength with Social Media. In: Proceedings of CHI, Boston (2009)
8. Guy, I., Ronen, I., Raviv, A.: Personalized Activity Streams: Sifting through the ”River of News”. In: Proceedings of RecSys, Chicago (2011)
9. Kaplan, A.M., Haenlein, M.: Users of the World, Unite! The Challenges and Opportunities of Social Media. Business Horizons 53(1), 59–68 (2010)
10. Muller, M.J., Freyne, J., Dugan, C., Millen, D.R., Thom-Santelli, J.: Return On Contribution (ROC): A Metric for Enterprise Social Software. In: Proceedings of ECSCW, Vienna (2009)
11. Noakes, M., Clifton, P.: The CSIRO Total Wellbeing Diet. Penguin Publ. (2005)
12. 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)
13. Shani, G., Gunawardana, A.: Evaluating Recommendation Systems. In: Ricci, F., Rokach, L., Shapira, B., Kantor, P.B. (eds.) Recommender Systems Handbook. Springer (2011)
14. 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)