Personalized Techniques for Lifestyle Change

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

Tasmanian ICT Centre, CSIRO
GPO Box 1538, Hobart, 7001, Australia
{jill.freyne,shlomo.berkovsky,gregory.smith}@csiro.au,
nbaghaei@unitec.ac.nz, stephenkimani@googlemail.com

Abstract. Online delivery of lifestyle intervention programs offers the potential to cost effectively reach large cohorts of users with various information and dietary needs. Unfortunately, online systems can fail to engage users in the long term, affecting their ability to sustain positive lifestyle change. In this work we present the initial analysis of a large scale application study of personalized technologies for lifestyle change. We evaluate the stickiness of an eHealth portal which provides individuals with three personalized tools – meal planner, social network feeds, and social comparison – to make change a reality in their lives. More than 5000 Australians took part in a 12 week study and provided solid empirical evidence for how the inclusion of personalized tools can assist and motivate users. Initial results show that the personalized tools boost user interaction with the portal, simplify information access, and assist in motivating users.

1 Introduction

The World Health Organisation is predicting that the number of obese adults worldwide will reach 2.3 billion by 2015 and the issue is attracting increased attention [13]. Much of this attention is being paid to online diet and lifestyle monitoring systems, which have been replacing traditional pen-and-paper programs. These systems include informative content and services, which persuade users to alter their lifestyle as well as tools and features which allow users to plan and record their progress. By the nature of these planning and recording tools, they gather a vast amount of information pertaining to dietary and exercise preferences of users. We propose harnessing this valuable user preference data to personalize the provided interactive features in order to reduce the work load of the individual and increase engagement with the system, and, in turn, the chances of sustained lifestyle change.

To investigate the role of personalized technologies in online lifestyle change systems, we designed and developed an experimental eHealth portal supported by an online social networking system and trialed it with a large cohort of users from across Australia. We implemented a number of intelligent personalized tools in the portal, some directly related to assisting users in their diet goals and some aimed at sustaining and increasing their interaction with the system. Specifically we added a *Personalized Meal Planner*, a *Personalized Network Activity Feed,* and a *Personalized Social Comparison* tool. The meal planner combines both explicit ratings on recipes

M. Peleg, N. Lavrač, and C. Combi (Eds.): AIME 2011, LNAI 6747, pp. 139–148, 2011.

and implicit feedback learned from a user's interaction with the planner to predict recipes liked by a user and produce meal plans relevant to a given day, based on food consumption patterns, frequency, and sequencing observed for that user. The network activity feeds highlight social network activities, which are highly relevant to a user, based on the observed relationship strength between the target user and the user who performed the activity and the user's interest in this type of activity. Finally, the social comparison tool also exploits user relationship strength and action relevance to select a set of highly relevant users and actions to ask social comparison questions about, in order to drive competition and positive feedback in the social network.

We conducted a large scale user evaluation of our lifestyle portal and the personalized tools. More than 5000 users participated in the study for a period of 12 weeks. We logged all interactions with the portal and analysed several parameters addressing the uptake of the personalized tools. Initial results show that the personalized tools boost user interaction with the portal, simplify information access, and assist in motivating users. Hence, the contributions of this paper are: 1) presentation of three personalized features for lifestyle change, 2) initial analysis of interaction with these features in a large scale live user study, and 3) report on learnings from the study on a range of topics from popularity of various features in the activity feeds to weight loss.

2 Related Work

Personalization and recommender technologies have been the proposed solution to the information overload for a number of years. It is well accepted that user modeling and the exploitation of user models to predict user needs and desires is an effective way of reducing the burden of users in domains such as information retrieval, e-commerce and entertainment. Here we touch on related work which employs personalization in the areas of meal planning, information access and social competition.

The use of implicit interaction and explicit rating data have both recently been explored in the area of food recommendations. Svensson *et al.* [10] report on their recommender system which judges the relevance of recipes based on the browsing patterns of users. Freyne *et al.* concentrate on explicit rating data on recipes and investigated three recommender strategies, which break down meals into ingredients to generating recommendations [2].

Social networking systems are continually changing and the challenge for individuals is keeping up with the actions and updates of their friends. Social Networking systems try to assist users by aggregating the actions of friends into Network Feeds which show in chronological order the activities of others. These feeds however do not consider the interests of the user or the relationship dynamics between friends on the system. Recently works have appeared which address the development of predictive models for computing the relevance of items within activity feeds. Gilbert and Karaholios developed a *tie strength* model [4], 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 evaluations were conducted with small cohorts of users. In contrast, 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 [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 a system for recommending SN activities of an interest based on long- and short-term models of content viewed and activities performed by users [14]. These systems were both evaluated using offline logs, whereas our work aims at live user evaluation.

Social comparison works by comparing the contribution of users to contribution of other users. Vassileva *et al.* used social visualization to increase participation by displaying the contribution made by each user and facilitating social comparison [11]. Harper et al. discovered that emails informing users whether their contributions was above or below average prompted users to rated more movies [5]. Michinov et al. showed the prolific impact of social comparison feedback on productivity of group members in an online collaborative brainstorming system [7]. To the best of our knowledge, no research addressed personalization aspects in social comparison.

3 Personalized Tools for Lifestyle Change

In our earlier work we investigated the role of social technologies for families interested in lifestyle change [1]. The study highlighted that the attitude towards health was correlated with engagement with the system. This prompted further investigation into increasing the value of the portal to individuals through personalized features, in order to prolong their interaction with the portal and, in turn, sustain the lifestyle change.

3.1 Meal Planner

Changing lifestyle primarily requires changing the types and amounts of food consumed and physical exercises performed. This is often a daunting task, as dietary habits are built over time and are hard to break. To combat this, some programs explicitly tell people what to eat or even supply the foods required. While this might be a short term solution, specified plans are often restrictive and may deter users. More importantly however, users do not acquire the diet management skills which they need to achieve long term success. The current alternative is to ask users to plan from scratch, which can be a daunting task. Our meal planner aims to assist people in planning desirable meal plans by acquiring a small amount of explicit meal preferences and learning from interactions of the users.

The domain of food is varied and presents a challenge for recommender systems. We gather explicit and implicit user preferences for recipes. We gather initial explicit recipe ratings on a 5-Likert scale, deliberately spanning a number of recipe categories to maximize the information gain. Also, we learn an implicit user profile with each meal planned by the user. We determine the implicit relevance of each recipe to the target user by examining how often the recipe appears in the user's meal plans. Finally, we combine the explicit ratings and implicit data in a weighted manner.

Our user profile is structured as a ratings vector, where each rating represents a recipe on which we have either explicit or implicit knowledge. We use a traditional collaborative filtering algorithm [6] to compute predictions for unrated recipes based on the ratings of N most similar neighbours. Briefly, neighbours are identified using Pearson's correlation algorithm shown in Equation 1 and predictions for recipes not rated by the target user are computed using Equation 2.

$$sim(u_a, u_b) = \frac{\sum_{i=1}^{k} \left(u_{ai_i} - \overline{u_a}\right)\left(u_{bi_i} - \overline{u_b}\right)}{\sqrt{\sum_{i=1}^{k} \left(u_{ai_i} - \overline{u_a}\right)^2} \sqrt{\sum \left(u_{bi_i} - \overline{u_b}\right)^2}} \tag{1}$$

$$pred(u_a, r_t) = \frac{\sum_{n \in N} sim(u_a, u_n) rat(u_n, r_t)}{\sum_{n \in N} sim(u_a, u_n)} \tag{2}$$

We employ a decay strategy for each recipe which takes into consideration how often a recipe appears in the user's meal plan and when it last occurred. Equation 3 shows the exponential decay algorithm which determines the score for a target recipe r_t for user u_a on day D where k<0. The final processing of the recommendation list occurs as the user interacts with the planner, such that recommendations that would break the diet rules based on the current partial plan disappear from the recommendation list.

$$score(u_a, r_t, D) = pred(u_a, r_t) * e^{kD} \tag{3}$$

Figure 1 illustrates the interface of the meal planner. The daily plan is shown in the centre, a structured tree of recipes is on the left, and the recommended recipes are on the right. Users can drag their preferred recipes to/from the daily plan and the recommended list changes accordingly.

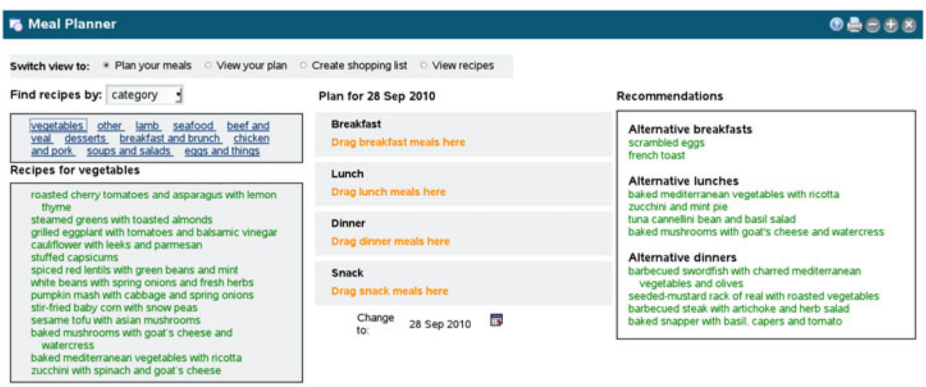

Fig. 1. Meal planner interface

3.2 Network Activity Feeds

One of the burdens on social networking systems (SNS) is their success. Typically SNS's are used for communication and sharing but billions of actions are carried out daily making keeping up to date extremely difficult. Although network feeds aggregate the activities of friends and deliver updates, they normally disregard the user's interests and the relevance of the feed content, leaving the viewer to search for interesting updates. Our personalized network feeds aim to reduce the burden of identifying interesting activities by computing the relevance of each activity and ordering the feed, such that the relevant activities are presented higher in the feed.

Fig. 2. Activity feed interface

We propose scoring feed activities based on previously observed interactions of the target user T with the social network. In short, the relevance score of a feed item (e.g., "*Bob added a photo*" in *Alice*'s feed), is computed as a weighted combination of the user-to-user score between *Alice* and *Bob* and user-to-action score between *Alice* and *adding photos*. Hence, the overall relevance score $S(T,I)$ is computed as a weighted combination of the user-to-user $S_U(T,u_x)$ and user-to-action relevance score $S_A(T,a_z)$:

$$S(T,I)=w_1 S_U(T,u_x)+w_2 S_A(T,a_z) \tag{4}$$

where w_1 and w_2 denote the relative weights of the components. In our case, we assign more weight to w_1, in order to emphasise activities performed by relevant users.

To compute the user-to-user relevance score $S_U(T,u_x)$, we adopt the model of [14] and use four categories of factors: (1) user factors (UF) – online behaviour and activity of the target user, (2) subject user factors (SUF) – online behaviour and activity of the subject user, (3) direct interaction factors (DIF) – direct communication between the two users, and (4) mutual connection factors (MCF) – communication between the users and their common network friends. Overall user-to-user relevance score $S_U(T,u_x)$ is computed as a weighted combination of the category scores:

$$S_U(T,u_x)=w_3 S_{UF}(T,u_x)+w_4 S_{SUF}(T,u_x)+w_5 S_{DIF}(T,u_x)+w_6 S_{MCF}(T,u_x) \tag{5}$$

In our case, we assign more weight to w_{DIF}, in order to emphasise the importance of direct communication between the users. In turn, 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 individual scores of observable network interaction factors in each category. Overall, we use 32 factors for the UF/SUF categories and 28 factors for the DIF/MCF categories.

The frequency of performing actions is the main indicator of user-to-action relevance scoring. 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 $S_A(T,a_z)$ is computed as the relative relevance of a_z for T and normalised by the relevance of a_z for all users:

$$S_A(T,a_z) = \frac{f(T,a_z)}{f(T)} / \frac{f(a_z)}{f()} \tag{6}$$

Figure 2 illustrates the interface of the activity feeds. 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. Users have the facility to adjust the number of items shown and seek for further items of interest.

3.3 Social Comparison

Social comparison, as a persuasive technique to affect user behavior and increase their engagement, has recently received much attention. Social comparison works by asking users to compare their friends in relation to a specific question, which can relate to either a user's online or offline activity. For example, "*Who spends more time on the site, Alice or Bob*?" or "*Who is more outgoing, Alice or Bob?*". The selected user (Alice or Bob) then receives feedback on this from the system. Although social comparison has been shown as an effective technique in changing online behavior and increasing user engagement, the reach and uptake of social comparison could be further improved through personalization. We implemented a personalized social comparison tool, which selects the topics and users which are best suited for the target user to compare, in order to maximize the chances of uptake and the generation of positive feedback from the system.

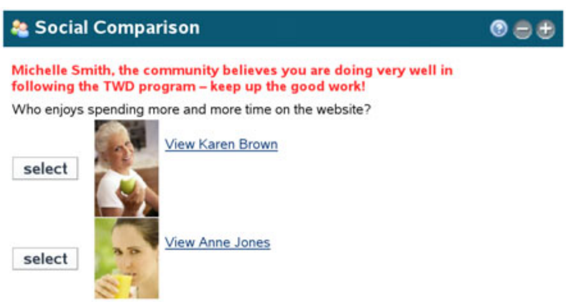

Fig. 3. Social comparison interface

We re-use the user-to-user $S_U(T,u_x)$ and user-to-action $S_A(T,a_z)$ relevance scores detailed above to personalize the questions and subjects that are presented to an individual. Our personalized social comparison tool aims to select candidates for comparison who meet the following criteria: (1) high user-to-user relevance $S_U(T,u_x)$, (2) high user-to-action relevance, and (3) low number of received feedback $fb(u_x)$ from system. The personalized social comparison initially selects M users ($u_1,\ldots,u_M$) with the highest user-to-user score $S_U(T,u_x)$. In order to prioritize users who received less feedback from the system, the relevance score of each of them is adjusted as in Equation 7 and two users with the highest prioritized score $S'_U(T,u_x)$ are selected as the candidate users CU_1, CU_2.

$$S'_U(T,u_x)=fb(u_x)(1-S_U(T,u_x)) \quad (7)$$

Then, we select candidate questions which meet the following criteria: (1) high user-to-action score $S_A(T,a_z)$, (2) high user-to-action scores $S_A(CU_1,a_z)$ or $S_A(CU_2,a_z)$ for either CU_1, or CU_2, and (3) not asked in previous K days, $lastask(T,q_x)$>K. We combine the user-to-action score of the target user and two candidate users as follows:

$$S_A(T,CU_1,CU_2,a_z)=S_A(T,a_z)S_A(CU_1,a_z)+ S_A(T,a_z)S_A(CU_2,a_z) \quad (8)$$

The action with the highest score is chosen as the candidate action to ask about. Each action is associated with a set of questions. If the question associated with the candidate action was not asked in previous K days, it is selected and asked. Figure 3 illustrates the interface of the social comparison.

4 Evaluation and Analyses

The eHealth portal we developed is a diet compliance system, providing users with skills, information, and tools designed to sustain and enhance their interaction with the portal in order to assist them with diet compliance and lifestyle chance. The core dietetic information was extracted from the CSIRO's Total Wellbeing Diet [8]. We evaluated the three personalized tools in a live study involving the users of the portal. More than 5000 users participated in the study for a period of 12 weeks, from September to November of 2010. Interaction with website features was optional, thus the number of users who interacted with each tool varied, we report here on the subsets of users who interacted with the features detailed above. The users were uniformly divided into several experimental groups, such that half of the groups were exposed to the personalized and half to non-personalized tools. As such, half of the users used the personalized and half the standard planner with no recommendations; half were exposed to personalize and half to non-personalized activity feeds; and lastly half were asked personalized and half randomized social comparison questions. No personalization was applied during the first week, due to the bootstrapping phase of the recipe preferences and relevance scores.

4.1 Meal Planner

Users in the personalized groups received recommendations on a daily basis through the meal planner. The recommendations were limited to three recommendations for *breakfast*, *lunch*, and *dinner* and were determined using the algorithms presented in Section 3.1. The recipes with the highest weighted relevance to the user were deemed recommendation candidates. Different users planned with different levels of granularity, i.e., for individual meals, a set of meals, or even for multiple days. For the sake of simplicity, we report on planning at the level of meals and days with plans.

Note that users who received personalized recommendations did not plan for a larger period of days than those who had no assistance. However, we do note that the plans created by those with assistance were more detailed with an average of 4.93

entries per day (including snacks) in comparison to 4.42 when no assistance was offered. Again, the total number of user interactions with the planner was comparable, as well as the number of required user interactions per meal planned.

Table 1. Uptake of the meal planner

	users	days planned	meals per day	interactions	interaction per meal planned
non-personalized	*512*	*9.17*	*4.42*	*100.69*	*2.11*
personalized	*515*	*9.11*	*4.99*	*102.48*	*2.13*

The meal planner system was designed to provide recommendations from a set of recipes associated with the diet program. We assumed that users would base their meal plans primarily around these recipes but in fact users primarily planned around recipes that they manually added to the planner. The first change requested by users was the facility to add their own recipes to the planner. Many users added their own recipes, others added combinations of foods (e.g., work lunch), and some simply worked with individual food groups (breads, proteins, etc). Overall, more than 12,000 extra meals were inserted to the plans. Thus, about 80% of all meals planned were from these additional items and our recommender could potentially only be effective in 20% of cases. Further to this, the interface did not allow users to easily browse recipes in the recommendation panel which was likely to have impacted the uptake.

4.2 Network Activity Feeds

Users in the personalized groups were exposed to feeds, in which the relevance scores were computed as presented in Section 3.2, and the activities with the highest score were presented high in the feed. Users in the non-personalized groups were exposed to non-personalized feeds, which presented the activities in reverse chronological order. The feeds were generated upon a user's login to the portal, such that the predicted scores of the activities were not re-computed until the next login. Table 2 summarises the number of users, sessions with feed clicks, clicks observed, and two click-through rates (CTR_u – number of clicks per user and CTR_s – number of clicks per session with clicks), computed 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.

Table 2. Uptake of the feeds

	users	sessions	clicks	CTR_u	CTR_s
personalized	*1397*	*390*	*901*	*0.6450*	*2.3103*
non-personalized	*1416*	*382*	*805*	*0.5685*	*2.1073*

Generally, the uptake of the activity feeds was not as high as we expected. There are several possible explanations to this. Firstly, unlike in other online social networks, users of our portal had no offline familiarity with each other. As a result, the friending level was low and the establishment of strong user-to-user relationships took longer. Secondly, many users requested to include a thumbnail image of users in the activity feeds. Supposedly, this could have increased the attractiveness of the feed items and boost the uptake of the feeds. Finally, about half of the feed clicks were observed for the first week of the study, for which we ignored the observed users' interaction due to the required bootstrapping of the user-to-user and user-to-action scores. This is in line with previous works on social network, which observe the highest drop-off rates at the initial stages of interaction [1].

4.3 Social Comparison

The personalized social comparison feature was the least useful feature of the portal, resulting in low response rates to the questions posed. Overall, we obtained only 315 questions received responses from a cohort of 44 users, who answered on average 7.15 questions each (stdev 19.7). The obtained feedback (i.e., users that were selected in the answers) addressed 254 users, averaging at 1.2 messages per user (stdev 1.09).

The social comparison feature was not well received by the users, with many of them feeling that the system was asking them to pass judgement on other users, which they were often uncomfortable with. The lower acceptance of social comparison tools in general was also observed in [12], which reports only a 5% update by users in traditional social networking systems. We observed a 7.2% initial uptake rate, but very few users re-used the feature afterwards. 1349 users were given the opportunity to engage with the social comparison tools, however only 44 took up the opportunity. Generally interaction with the feature was low but one user answered every question (over 100) which was generated for him. Due to the low uptake, it is difficult to asses the impact of the personalization, but there seem to be differences in uptake of questions: users shown non personalized questions responded to an average of 2.53 questions and users shown personalized questions responded to on average 2.89 questions. A further separate analysis of this tool in a different environment is required.

The social comparison feature did not have the desired impact on users and, in fact, it had quite an opposite impact. Some users were unsure as to how and why they should respond to the questions. Many of them were uncomfortable with comparing other users with respect to their performance on the portal. Again, several users complained that the social comparison interface was cumbersome and inconvenient. Since social comparison was not directly related to the tasks associated with the diet, users felt that it was unproductive for them.

5 Conclusions and Future Work

In this paper we report on a large scale live user study on the impact of personalization in an online portal for lifestyle change. We examine the update of three personalized tools with differing roles in the portal. Our meal planning tool was central to the portal and diet, the activity feeds were central to maintaining awareness and interaction with the portal, and the social comparison was an experimental tool aimed to motivate users. More than 5000 users from across Australia participated in the study for a period of 12 weeks. We have presented in the paper our initial findings, which suggest that personalized tools have the ability to boost user interaction, simplify information access, and motivate users. We also presented our observations and feedback obtained from users on the personalized tools and interaction with them.

Next steps for this work are obviously to carry out a more in-depth analysis of the overall impact of all of the features evaluated in this study. In this work, we have looked at each feature individually, whereas the larger picture can only be known by looking at the portal in its entirety. Thus, we will analyse how the combination of the provided tools impacted weight loss, attitude of users, duration and intensity of interaction, and engagement with the portal and lifestyle change program.

Acknowledgements

This research is funded by the Australian Government through the Intelligent Island Program and CSIRO Food and Nutritional Sciences. The Intelligent Island Program is administered by the Tasmanian Department of Economic Development, Tourism, and the Arts. The authors acknowledge Mealopedia.com and Penguin Group for permission to use their data.

References

1. Baghaei, N., Kimani, S., Freyne, J., Brindal, E., Smith, G., Berkovsky, S.: Engaging Families in Lifestyle Changes through Social Networking. International Journal of Human-Computer Interaction (2011) (in press)
2. Freyne, J., Berkovsky, S.: Recommending Food: Reasoning on Recipes and Ingredients. In: Proceedings of UMAP (2010)
3. Freyne, J., Berkovsky, S., Daly, E.M., Geyer, W.: Social Networking Feeds: Recommending Items of Interest. In: Proceedings of RecSys (2010)
4. Gilbert, E., Karahalios, K.: Predicting Tie Strength with Social Media. In: Proceedings of CHI (2009)
5. Harper, F.M., Li, S., Chen, Y., Konstan, J.: Social Comparisons to Motivate Contributions to an Online Community. In: de Kort, Y.A.W., IJsselsteijn, W.A., Midden, C., Eggen, B., Fogg, B.J. (eds.) PERSUASIVE 2007. LNCS, vol. 4744, pp. 148–159. Springer, Heidelberg (2007)
6. Konstan, J., Miller, B., Maltz, D., Herlocker, J., Gordon, L., Riedl, J.: GroupLens: applying collaborative filtering to Usenet news. Communications of the ACM 40(3) (1997)
7. Michinov, N., Primois, C.: Improving productivity and creativity in online groups through social comparison process: new evidence for asynchronous electronic brainstorming. Computers in Human Behavior 21(1) (2005)
8. Noakes, M., Clifton, P.: The CSIRO Total Wellbeing Diet. Penguin Publ., Australia (2005)
9. van Pinxteren, Y., Geleijnse, G., Kamsteeg, P.: Deriving a recipe similarity measure for recommending healthful meals. In: Proceedings of IUI (2011)
10. Svensson, M., Hook, K., Laaksolahti, J., Waern, A.: Social navigation of food recipes. In: Proceedings of CHI (2001)
11. Vassileva, J., Sun, L.: An improved design and a case study of a social visualization encouraging participation in online communities. In: Haake, J.M., Ochoa, S.F., Cechich, A. (eds.) CRIWG 2007. LNCS, vol. 4715, pp. 72–86. Springer, Heidelberg (2007)
12. Weiksner, G.M., Fogg, B.J., Liu, X.: Six patterns for persuasion in online social networks. In: Oinas-Kukkonen, H., Hasle, P., Harjumaa, M., Segerståhl, K., Øhrstrøm, P. (eds.) PERSUASIVE 2008. LNCS, vol. 5033, pp. 151–163. Springer, Heidelberg (2008)
13. World Health Organisation Chronic disease information sheet, `http://www.who.int/mediacentre/factsheets/fs311/en/index.html` (accessed January 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 (2010)