Automatic User Adaptation for Behavior Change Support*

Ulrich Reimer
Institute for Information & Process Management
University of Applied Sciences
St. Gallen, Switzerland
ulrich.reimer@fhsg.ch

Edith Maier
Institute for Information & Process Management
University of Applied Sciences
St. Gallen, Switzerland
edith.maier@fhsg.ch

Tom Ulmer
Institute for Information & Process Management
University of Applied Sciences
St. Gallen, Switzerland
tom.ulmer@fhsg.ch

ABSTRACT

User adaptation is crucial to account for the heterogeneity of target user groups and thus to increase the effectiveness of a system aimed at encouraging behavioral change. The paper first introduces an application framework that comprises various components which allow to accommodate user preferences. It then describes in detail the algorithms for user adaptation. To this end, approaches from user modeling as well as from collaborative filtering are employed. The result is a self-learning application that changes in line with a user's progress, which is expected to enhance user acceptance as well as increase and sustain the motivation for behavioral change. The framework is currently being implemented in a mobile health app and will subsequently be evaluated.

Keywords

Behavioral Change Support, Mobile Health, User Adaptation, User Modeling, Collaborative Filtering, Self-Learning

1. INTRODUCTION

Behavioral economics is an approach that promises to ameliorate the shortcomings of traditional healthcare management, especially with regard to chronic disease (see e.g. [3]). Behavioral economists use knowledge from behavioral science as well as motivational psychology and neuroscience to study how individuals make decisions which are often non-rational, and biased by a series of mental shortcuts, for instance, the so-called "status quo bias" [7]. Apart from the status quo bias, people's behavior is also susceptible to the influence of default rules, framing effects and starting points. Consequently, persuasion strategies can involve changing the way options are presented, e.g. by adapting the rules that drive user interaction in a mobile health app.

The philosophy of behavioral economics is also called "libertarian paternalism", namely that people should not be forced to act in certain ways, but rather encouraged to act in ways that are better for them or help them stopping bad habits formed over time. This idea of a "nudge" favors invitations to change behavior by means of brief persuasive interventions rather than by constraints and sanctions [20].

Most mobile health solutions, i.e. mobile devices connected to medical applications or sensors, as well as pure lifestyle apps actually include some kind of support for the users to achieve their goals. However, these nudges tend to be hardwired, i.e. they do not adapt to user preferences and needs and on the whole they are not grounded in behavioral change theory (see e.g. [15]). A comprehensive and coherent framework of behavior change interventions has been developed by Susan Michie and her co-workers [11]. The framework, coined "behaviour change wheel", distinguishes between different kinds of interventions, among them persuasion, incentivization and coercion. These intervention types are also addressed by behavior change support systems (BCSS) as introduced and defined by [12]:

"Behavior change support systems (BCSS) are information systems designed to form, alter, or reinforce attitudes or behaviors or both without using coercion or deception."

The persuasive systems design (PSD) model, a framework for designing a BCSS introduced in [13], draws from the seminal work by Fogg on persuasive technology [4].

The PSD model has been elaborated on by [14] who have created a conceptual framework of what they call *tailoring*. It includes seven inter-related key concepts: feedback, interhuman interaction, adaptation, user targeting, goal setting, content awareness, and self-learning. According to the authors, a self-learning application is able to update its internal model of the user by recording and learning from the interactions the user has with the application and thus to change with the user over time. Self-learning corresponds to the automatic user adaptation as employed in our BCCS framework.

In their survey of physical activity monitoring and tailoring, [14] found that most applications only employ feedback – the most obvious form of tailoring. They found only one paper [2], which describes a platform designed to stimulate healthy living using more than two different tailoring concepts. The paper, however, is primarily conceptual in nature and does not provide any detail on technical implementation or algorithms nor does it refer to any behavior change theories [2].

If we include areas besides physical activity, we do encounter other applications that incorporate different tailoring approaches, for example, the food recommender system developed by [5], which not only offers recipe recommendations in accordance with user's preferences but also takes into account their health status.

In this paper, we propose a self-learning application which also uses other tailoring techniques as defined by [14] such as context awareness, goal setting and user targeting.

*The research presented in this paper has been made possible by a grant from the Swiss Gebert Rüf Foundation. We also like to thank the anonymous reviewers of an earlier version for their helpful comments.

2. MAIN CONSTRUCTS FOR BEHAVIORAL CHANGE SUPPORT

Our work on user adaptation is based on a BCSS application framework which is described in more detail in [17]. In the following, we will shortly introduce its main constructs.

2.1 Goal Hierarchies

At the heart of any BCSS, esp. for mobile health, are the goals a user wants to achieve. Our BCSS application framework [17] therefore includes a construct for specifying one or more *goal hierarchies* for the targeted application domain of a BCSS. A goal hierarchy starts with a top goal which represents a user's primary goal. The top goal tends to be *long-term*, it may be measurable, e.g. "body mass index below 30 within 6 months", or be more generic, e.g. "keep healthy". It can usually be achieved in a variety of different ways, e.g. by engaging in physical activity, lowering the stress level, by eating regular meals or a combination thereof. Each option is represented by a sub-goal, together with an indication if the sub-goal is sufficient for reaching the higher-level goal or if several sub-goals need to be reached. Sub-goals can be broken down into further sub-goals until these can be associated with a measurable activity. We call such goals operationalized:

Definition:
An *operationalized goal* is short-term and is associated with a measurable activity to reach the goal.

2.2 Nudge Types

In the course of our research we have conducted extensive interviews with potential end-users which confirm the findings of other researchers [6, 8, 16], namely that behavior is influenced by a variety of factors, e.g. age, sex, socio-economic status, attitudes, personality, social environment and peer group. Most existing BCSS, however, have implemented only a fixed or limited set of interventions (nudges) that do not take into account the *heterogeneity of target users*. This results in low intervention effectiveness and poor user acceptance. Our framework incorporates a variety of *nudge types* such as suggestions, praise and rewards, which are selected on the basis of goal achievement graphs.

2.3 Goal Achievement Graphs

We control the triggering of nudges by using *goal achievement graphs* (GA graphs) that reflect a user's progress by plotting goal achievement along the time axis. GA graphs make sense for goals which are reached by repeated single activities, such as 10,000 steps per day or self-weighing three times a week. If a goal can be reached with a single activity (e.g. self-weighing once per day), reminders are sufficient, which means that the possibility to support users in achieving their goals is limited.

The rectangular area spanned by a GA graph is divided into three sub-areas (see Fig.1): Area A signifies good progress, area B indicates slow progress, while area C indicates that the user might not reach the goal. There is one GA graph for each operationalized sub-goal.

A variety of trigger rule schemas are predefined [19], e.g. to generate a praise when the user is performing well and progress lies in area A. If progress lies within area B, nudges such as suggestions or reminders are triggered. When goal achievement falls into area C, stronger nudges may be called for because the user is at risk of missing the goal. On the other hand, if the user catches up and moves back into area B or from area B to A, a praise may be generated.

The boundaries between the three sub-areas change in line with a user's typical progress with regard to his or her goal achievement activities (see Sec.3.1.2). Figure 1 shows an initial setup.

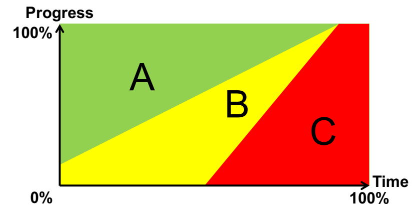

Figure 1: Initial setup of a goal achievement graph.

3. USER ADAPTATION

Personalizing a BCSS to a user's individual needs and preferences requires some effort from the user. As shown by our interviews the average user will shun this additional effort and find it difficult deciding which choices to make, e.g. between alternative ways to reach a goal for physical activity or which kinds of nudges to prefer. We therefore aim at enabling the system to automatically adapt to the user. Depending on the target of adaptation either user modeling [9] or collaborative filtering [1] may be more suitable. In the former case, it is based on the user's behavioral history and is independent from other users of the system (see Sec.3.1). In the latter case, adaptation draws on the preferences of similar users of the same system and requires a sufficiently large number of users with a sufficiently long history (see Sec.3.2).

3.1 User Modelling

User modeling implies the successive build-up of user models based on user behavior and from that derive how the system should interact with the user. Adaptation through user modeling is applied to two constructs in our BCSS framework: nudge types (Sec.3.1.1) and GA graphs (Sec.3.1.2).

3.1.1 Adapting Nudge Types to Users

Since it has been shown that the intervention of a persuasive system needs to be tailored to a specific user [8] we devised an algorithm to automatically adapt preferred nudge types to users. For example, if a user repeatedly follows a suggestion made by the system, this is a good indicator that the user responds well to suggestion nudges. Also, whilst some users might respond well to reminders or feedbacks that they are falling behind their peer group, other users might simply ignore such messages.

The adaptation is based on which kinds of nudges have proved more successful than others. The system starts by selecting nudge types randomly and monitoring how well each of them works. Once the system has identified the nudge types that are more successful in terms of triggering intended behavior, the system uses those types more often.

Algorithm 1 describes the underlying principles of our approach in more detail. The following remarks apply:

Line 1: The algorithm is running in the background and gets active whenever a nudge is triggered.

Line 2: We are not interested in individual nudges but in nudge types so that all further calculations refer to the type of nudge just triggered.

Lines 3–4: t_1 and t_2 are relative to the time interval during which the goal is to be achieved. For example, if this time interval is a day, then t_1 and t_2 would be times of day. In case of a weekly time interval t_1 and t_2 would be the day of the week. Note that for the subsequent nudge, t_2 becomes t_1 in the next iteration of the for-loop.

Line 5: Within a goal achievement interval such as a day or a week, the function $p(t)$ yields the mean progress at time t the user has achieved in the course of his or her behavioral history, e.g. for daily goals the mean progress at t='noon' over all the days goal achievement has been monitored.

Line 6: The value of the variable $meanProgress$ is the user's mean progress increment made so far over all times. We calculate this by taking the differential of p over the time variable t and sum up its values (i.e. the ascent of p) at all time points in the time interval within which the goal is to be achieved, i.e. between t_{start} which stands for the beginning of the time interval and t_{end} which stands for its end. The increments between the time points are not specified, but could e.g. be minutes if the time interval is a day. We normalize the sum by dividing it by the number of time increments in the time interval.

Line 7: The progress the user has made towards goal achievement after the nudge n, is calculated and divided by the time duration between the two nudges. The time units of t_1 and t_2 must be the same as in Line 6, e.g. minutes. Thus the resulting relative progress $relProgr$ is independent of any particular time interval.

Lines 8–10: The mean progress the user has made after a nudge of type $nType$ gets updated.

Line 11: The impact score of nudges of type $nType$ is the ratio of the mean progress after that particular nudge type and the overall mean progress made by the user.

The nudge types are ranked according to their impact scores. Those with the highest scores are used more often since they (appear to) work better for the particular user. The scores with a value below 1 are used less often. They are not completely blocked because they might work better in the future in which case their scores can increase again.

3.1.2 Adapting Goal Achievement Graphs

The boundaries between the three areas in the GA graphs (cf. Fig.1) are constantly adapted to reflect the user's typical timing of when he or she performs the activities required to reach the corresponding goal. To this end, the system aggregates the progress curves of all successful goal achievements into one curve. This is done by taking the mean of all related achievement values from the past for each point on the time axis, which results in a typical goal achievement curve for the user. We take this curve as the boundary between areas A and B so that as long as the user is above the mean progress curve he or she stays in area A. The boundary between areas B and C is set to the mean of the A-B boundary and the slowest overall progress curve the user has ever had whilst still achieving the goal.

To reflect a user's change of habits, the more recent a goal achievement curve, the higher the weight it is given when calculating the aggregated curve. In this way, the GA graph reflects habit changes over time (for more details cf. [19]).

Algorithm 1 Continuous update of the impact score for each nudge type

1: **for all** triggered nudges n **do**
2: $nType \leftarrow type(n)$
3: $t_1 \leftarrow$ time of the nudge n
4: $t_2 \leftarrow$ time of the next nudge or of goal achievement (whichever applies)
5: update the function p – the aggregated mean of all progress curves in the user history – for the new time segment $[t_1, t_2]$
6: $meanProgress \leftarrow \frac{1}{|[t_{start}, t_{end}]|} \sum_{t_i \in [t_{start}, t_{end}]} \frac{\mathrm{d}p}{\mathrm{d}t}(t_i)$
7: $relProgr \leftarrow \frac{achievement(t_2) - achievement(t_1)}{t_2 - t_1}$
8: $accRelProgr_{nType} \leftarrow accRelProgr_{nType} + relProgr$
9: $count_{nType} \leftarrow count_{nType} + 1$
10: $meanProgr_{nType} \leftarrow accRelProgr_{nType} / count_{nType}$
11: $score_{nType} \leftarrow meanProgr_{nType} / meanProgress$
12: **end for**

3.2 Collaborative Filtering

Our BCSS framework makes use of *collaborative filtering* [1] to generate recommendations concerning goal setting, i.e. which sub-goals to pursue to reach a main goal and which target values to adopt. To utilize collaborative filtering we first need to define the features with which to describe a user and a similarity measure on those features. There are two principle approaches for characterizing a user: first, by means of the user's socio-demographic data, and secondly, by his or her past behavior. Clearly, in the case of a BCSS the user's behavior is significantly more relevant than the socio-demographic data. For example, people with different socio-demographic background are likely to behave in a similar way when it comes to maintaining weight loss.

We therefore decided to characterize users by their activity profile. To this end, the BCSS records the user's activities using the sensors connected to the system. For example, with accelerometers it is rather straightforward to detect basic activities such as standing, sitting, walking, running, cycling [10]. The collected activity history is mapped into a *feature vector* which is then used for describing a specific user. For each activity the feature vector includes a value which gives the average percentage of time per day which the user has spent on a particular activity over all past days: $\langle t_{avg}(sit), t_{avg}(stand), t_{avg}(walk), t_{avg}(run), t_{avg}(cycle), \ldots \rangle$

Additionally, we add the user's stress profile (using our SmartCoping system [18]) to the feature vector which consists of the average time during a day with high, medium and low stress:
$\langle t_{avg}(high\text{-}stress), t_{avg}(medium\text{-}stress), t_{avg}(low\text{-}stress) \rangle$
Similarity between two users is measured using the *cosine* on their feature vectors. Recommendations are generated only from users with a sufficiently long usage history. The recommendations generated for a user u concern:

a) Sub-goals: All similar users with the same main goals as user u are considered. For each main goal, sub-goals are suggested with which to achieve the main goal. To this end, a score is calculated for each sub-goal sg. The sub-goals with the highest scores are recommended.

$$score(sg) = \frac{1}{1+n_{sg}} \sum_{i=1}^{n} sim(u,u_i) \cdot has\text{-}subgoal(u_i, sg)$$

where

$$has\text{-}subgoal(u,g) = \begin{cases} 1\text{: if user } u \text{ has subgoal } g \\ 0\text{: } else \end{cases}$$

$$n_{sg} = \sum_{i=1}^{n} has\text{-}subgoal(u_i, sg)$$

and $sim(u,u')$ = cosine over the feature vectors associated with users u and u'.

b) Target values: For all recommended sub-goals sg as well as for sub-goals already set, a target value is suggested:

$$goal\text{-}value(sg,u) = \frac{1}{n_{sg}} \sum_{i=1}^{n} sim(u,u_i) \cdot goal\text{-}value(sg,u_i)$$

A possible extension would be to include the *success rate of similar users* with their goal settings, to recommend not only what other users do, but what has been most successful.

4. CONCLUSIONS AND OUTLOOK

We have introduced various constructs of an application framework for goal setting and generating nudges to encourage behavioral change. To do justice to the heterogeneity of target users, our system automatically adapts to a user's needs and preferences as well as his or her changing behavior. For this purpose, we presented user modeling approaches that take into account a user's behavioral history as well as collaborative filtering techniques that draw on the collected evidence from other users of the system. Together, they enable user-specific nudges as well as recommendations concerning setting one's goals and how best to achieve them.

The adaptation algorithms are being implemented in a mobile health app which will subsequently be evaluated. In the future we might include additional features such as heart rate in combination with accelerometers to identify activity intensity.

We also plan to extend our framework to cover other interventions for inducing behavioral change, such as incentivization (rewards) and environmental restructuring (social support) (cf. [11]).

5. REFERENCES

[1] G. Adomavicius and A. Tuzhilin. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17(6):734–749, 2005.

[2] P. Bielik, M. Tomlein, P. Krátky, Š. Mitrík, M. Barla, and M. Bieliková. Move2play: an innovative approach to encouraging people to be more physically active. In *Proceedings of the 2nd ACM SIGHIT international health informatics symposium*, pages 61–70, 2012.

[3] European Commission. *Green paper on mobile health.* 2014.

[4] B. J. Fogg. *Persuasive Technology: Using Computers to Change What We Think and Do.* Morgan Kaufmann, 2002.

[5] M. Ge, F. Ricci, and D. Massimo. Health-aware food recommender system. In *Proceedings of the 9th ACM Conference on Recommender Systems*, pages 333–334. ACM, 2015.

[6] S. Halko and J. A. Kientz. Personality and persuasive technology: An exploratory study on health-promoting mobile applications. In *Proceedings of the 5th International Conference on Persuasive Technology*, PERSUASIVE'10, pages 150–161, Berlin, Heidelberg, 2010. Springer-Verlag.

[7] D. Kahneman. *Thinking, fast and slow.* Penguin Books, New York, Toronto, London, 2011.

[8] M. Kaptein, J. Lacroix, and P. Saini. Individual differences in persuadability in the health promotion domain. In T. Ploug, P. F. V. Hasle, and H. Oinas-Kukkonen, editors, *PERSUASIVE*, volume 6137 of *Lecture Notes in Computer Science*, pages 94–105. Springer, 2010.

[9] A. Kobsa. Generic user modeling systems. *User Modeling and User-Adapted Interaction*, 11(1-2):49–63, 2001.

[10] J. R. Kwapisz, G. M. Weiss, and S. A. Moore. Activity recognition using cell phone accelerometers. *SIGKDD Explorations Newsletter*, 12(2):74–82, 2011.

[11] S. Michie, M. M. van Stralen, and R. West. The behaviour change wheel: A new method for characterising and designing behaviour change interventions. *Implementation Science*, 6(1):1–11, 2011.

[12] H. Oinas-Kukkonen. Behavior change support systems: The next frontier for Web Science. In *Proceedings of the Web Science Conference 2010*, 2010.

[13] H. Oinas-Kukkonen and M. Harjumaa. Persuasive systems design: Key issues, process model, and system features. *Communications of the Association for Information Systems*, 24(1):28, 2009.

[14] H. op den Akker, V. M. Jones, and H. J. Hermens. Tailoring real-time physical activity coaching systems: a literature survey and model. *User modeling and user-adapted interaction*, 24(5):351–392, 2014.

[15] M. Patel, D. Asch, and K. Volpp. Wearable devices as facilitators, not drivers, of health behavior change. *JAMA*, 313(5):459–460, 2015.

[16] S. Prost, J. Schrammel, K. Röderer, and M. Tscheligi. Contextualise! personalise! persuade!: A mobile HCI framework for behaviour change support systems. In *Proceedings of the 15th International Conference on Human-computer Interaction with Mobile Devices and Services*, MobileHCI '13, pages 510–515. ACM, 2013.

[17] U. Reimer and E. Maier. An application framework for personalised and adaptive behavioural change support systems. In *Proc. 2nd Int. Conf. on Information and Communication Technologies for Ageing Well and e- Selfealth (ICT4AWE)*. 2016.

[18] U. Reimer, E. Maier, E. Laurenzi, and T. Ulmer. Mobile stress recognition and relaxation support with SmartCoping: user adaptive interpretation of physiological stress parameters. In *Proc. Hawaii Int. Conf. on System Sciences HICSS-50*, 2017.

[19] U. Reimer, E. Maier, and T. Ulmer. A self-learning application framework for behavioral change support. *To be published*, 2016.

[20] R. H. Thaler and C. R. Sunstein. *Nudge: improving decisions about health, wealth, and happiness.* Penguin Books, New York, Toronto, London, 2009.