A Self-Learning Personalized Feedback Agent for Motivating Physical Activity

Harm op den Akker
Roessingh Research and
Development
Telemedicine Cluster
h.opdenakker@rrd.nl

Laura S. Moualed
University of Twente
Biomechanical Engineering
Group
l.s.moualed@student.utwente.nl

Valerie M. Jones
University of Twente
Telemedicine Group
v.m.jones@utwente.nl

Hermie J. Hermens
Roessingh Research and
Development
Telemedicine Cluster
h.hermens@rrd.nl

ABSTRACT
An important aspect in the treatment of various chronic diseases is to optimise physical activity levels. We present a general approach for the implementation of an electronic Feedback Agent that serves as a personal coach for achieving and maintaining a healthy level of physical activity through sustainable behavioural change. The Feedback Agent is a self-learning, context aware, personalized software agent that runs on the user's Smartphone and uses an external inertial sensor to keep track of the user's level of physical activity throughout the day. We highlight the three important aspects of feedback in our framework: the timing, content and representation of given feedback messages. Tailoring and optimization of feedback timing and content is in an advanced stage of research, while the representation aspect is largely a matter of future work.

Categories and Subject Descriptors
I.2.1 [**Applications and Expert Systems**]: Medicine and science; I.2.6 [**Learning**]: Induction

General Terms
Algorithm, Human Factors

1. INTRODUCTION
One of the focus areas in the treatment of chronic disease is to optimise physical activity levels. At Roessingh Research and Development we are investigating how to support this by objectively measuring a patient's activity levels throughout the day and providing regular feedback on the patient's performance via a Smartphone. The positive effects of regular feedback on behaviour change is supported by numerous scientific studies. For example, [5] showed the positive effect of regular feedback on four different kinds of health related behavioural change including increase of physical activity. Multiple studies using the Smartphone and activity sensor on patients with Chronic Low Back Pain at Roessingh Research and Development have shown the positive effects of giving feedback on the patient's physical activity throughout the day [12], while similar effects have been shown for patients with COPD and Chronic Fatigue Syndrome.

Although our general approach seems to be effective, as yet the underlying models and the methods applied to give feedback to the patients are not very sophisticated, being based on a "one size fits all" approach. We believe that, in order to be as effective as possible, the feedback methods should be adapted to the individual patient. This view is supported by a study by Noar et al., who in a meta analysis of 57 case studies, looked at the effects of indivually tailoreded, printed messages in health behaviour change. Tailored messages were found to be more effective than non-tailored messages [7]. Looking at the patient's personal preference, when asked, 80% showed a positive interest in the idea of receiving tailored information to help the self-management of their physical activity levels [6], bolstering our belief in the need for strong personalisation in our physical activity treatment.

In this paper we describe our general approach in giving tailored feedback to patients on their physical activity behaviour in order to balance or increase overall levels of activity. The aim is to design and develop an electronic Feedback Agent that runs on a Smartphone and functions as a personal activity coach to the patient, adapting itself to the patient's preferences and growing with the patient as they progress through the various stages of behavioural change [10].

The rest of this paper is outlined as follows: Section 2 gives a high level overview of how the Feedback Agent operates in different phases during use. Then, Section 3 provides details on the three important aspects of feedback: timing, content and representation. The paper concludes with a discussion on encountered and foreseen issues with the Feedback Agent in Section 4.
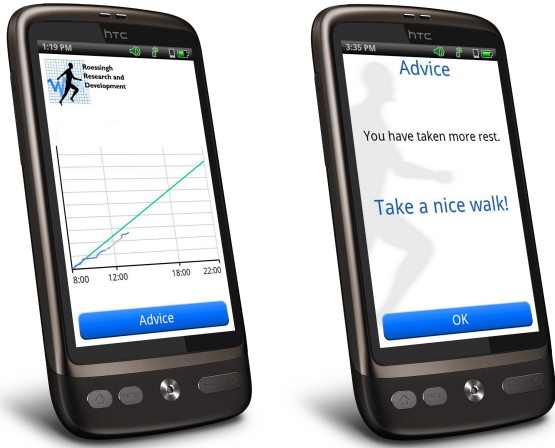
Figure 1: The interface for the Smartphone application that shows a graph of the user's current activity (left), and provides the user with textual feedback messages (right).

2. APPROACH

The proposed Feedback Agent is a Java/Android application running on a Smartphone that should be carried around by the user at all times. The user also carries a ProMove3D inertial sensor node[1], attached with a belt on the patient's hip, that constantly measures current activity intensity and transmits the measurements to the Smartphone over Bluetooth. The purpose of the Smartphone application is twofold: first it shows to the user their current activity, in a graph, together with a pre-defined "healthy reference" activity pattern, and second, it attempts to motivate the user to adhere to the baseline by sending text-based reminders when the user deviates too much from it. Figure 1 shows two screenshots of the current implementation of the feedback application, showing the graph of a user's current activity plotted over the reference line, and an example of a feedback message to encourage activity.

The tailoring of the feedback messages occurs in three steps:

1. The patient is invited to answer a short questionnaire about their current stage of change, personal situation (do you have a dog, a bike, a garden, etc...). The user is then asked to indicate their preferences to a set of sample feedback messages. Based on his answers, the user is automatically put into a cluster of people who have already been using the system and responded well to feedback messages for which the user has indicated preference. This constitutes the provisional starting user profile which in later phases is modified by application of a machine learning algorithm. The information is used in the process of generating the content of the feedback messages, as explained in more detail in Section 3.2. We assume that, based on the small amount of information regarding indicated preferences, we can succesfully cluster users, and thus al-

---

[1]http://www.inertia-technology.com

ready improve the feedback mechanisms compared to giving non-tailored feedback.

2. The Feedback Agent then starts operating in the *Cold Start Phase*. In this phase no "usage information" is available for the user yet, so the Feedback Agent has to decide on the timing, content and representation of feedback messages based on (similar) previous users of the system. At regular time intervals, the Feedback Agent predicts whether or not it is a good time to send a feedback message to the patient by querying the cold start classifier, which has been trained with data from previous research and other users that are currently using the system. Similarly, the content of the feedback message is generated based on statistical information used from other users predicted as similar to this user based on the clustering in step 1. During this phase, the responses of the user (whether or not they actually comply to a feedback message) are stored, together with all the relevant information regarding the context in which the messages was provided to the user.

3. When enough feedback responses from a particular user are stored, the Feedback Agent switches to its *Real Time Phase*. The stored responses to feedback messages are used to train a classifier for predicting the optimal timing of feedback messages, using genetic algorithms to decide on the optimal feature set to use [3]. This classifier is then used to predict at each fixed time interval whether or not it is a good time to give feedback (as explained in Section 3.1). The user's compliance to the feedback messages are also used to create an internal database that keeps track of how the user responded to messages that are phrased in a particular way. If for example, the user responded well to short messages, messages for which the *verbosity* dimension is set to *short* will be picked with a higher probability. After a feedback message is presented to the user, the application stores the time at which the user has actually seen the message. Then, after a certain time interval[2], the system can calculate whether or not the user complied to the previous message. If, for example, the message was meant to *encourage* more activity, the user *complied* if he became more active. This new feedback instance is then used to re-train the internal classifier to predict feedback timing, and the database for feedback content statistics is updated accordingly.

After the initial clustering and the cold start phase, the Feedback Agent will automatically keep updating its "knowledge" about the user, effectively tailoring the timing and content of feedback messages to the user. The information gathered from each individual user can also be sent back to a central server, where the data can help improve the quality of the cold-start classifiers.

3. THREE ASPECTS OF FEEDBACK

In our Feedback Agent framework, providing motivational feedback messages has three angles of approach (Figure 2). The *timing*, and *content* are two aspects that have already

---

[2]We currently use an arbitrary value of 30 minutes, which will be adapted if evidence suggests this value needs to be optimised.

been mentioned in the general approach overview in the previous section. A third aspect is the *represenation* of the message to the user. In this section, more detail is given about the work done on these three separate aspects of feedback.
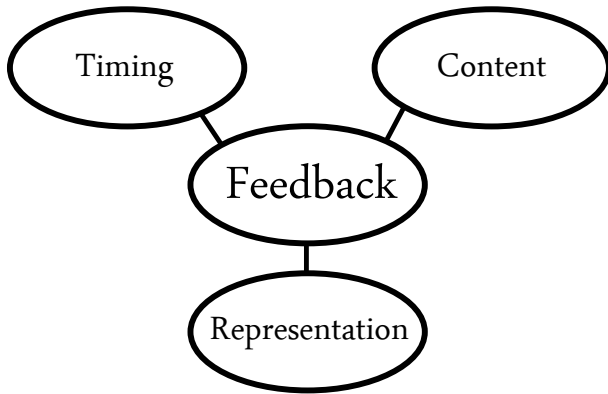

Figure 2: Feedback can be characterized in our framework by three aspects: the timing (when and how often to give feedback), the content (what information is conveyed through the message) and the representation (how is the information presented visualy/auditory).

3.1 Timing

When and how often to give feedback to a user is an important aspect when designing a remote monitoring and treatment application such as the described Feedback Agent, which is intended to encourage sustainable behavioral change. The user of the system should perceive the help they receive from the application as supportive and not as annoying. To minimize the *perceived interruption burden* is a research topic that is related to mobile devices in general. For example, a study by Ho and Intille shows that people are more receptive to notifications from their mobile device if the timing of the notifications coincides with activity transitions, compared to messages delivered at random times [4].

Instead of using a theory based model to infer appropriate moments for feedback, the approach we are taking is based on statistical inference. Earlier work by Op den Akker et al., shows the possibility of succesfully predicting compliance to a feedback message by looking at contextual features [9]. Feedback events from earlier studies on patients with Chronic Low Back Pain, Obesity and Chronic Fatigue Syndrome have been collected and augmented with contextual features such as weather information, history of use, and time related features. Using these feedback instances, a Ridor (RIpple-DOwn-Rule learner) classifier has been trained to succesfully predict compliance in 86% of the cases using Leave-One-Out evaluation. The results obtained in this study have been used to design and implement a software module for the Smartphone feedback application that can do on-the-fly compliance prediction for candidate feedback events [8].

The application keeps track of the user's current context parameters such as his activity level compared to the reference line, how long ago he received feedback, what his average compliance ratio is, etc. Also, a location module and weather module are in place to keep track of the current weather at the location of the patient, using the phone's network localization and the Google weather API. The gathered context information form the feature vectors for candidate feedback instances, which in the cold start phase are classified to predict compliance (the binary class attribute, as either yes or no). Thirty minutes after a given feedback message was seen by the user, the actual compliance to that instance is calculated by comparing the difference in activity, relative to the reference line at the time of the feedback message with the activity relative to the reference line at the current time. The actual compliance is stored with the feedback instance; ready to be used for the online classifier. If enough instances are gathered, the system will start training a new Ridor classifier. For this, the WEKA framework implementation [14] has been modified to run on Android devices, and extended with a genetic algorithm module to perform feature selection. The new *real time* classifier will then be used to predict compliance for newly generated candidate feedback events, and can be re-trained either daily or each time new user information (i.e. feedback instances for which the actual compliance is known) becomes available.

The software module described here for automatic selection of appropriate feedback timing has currently been implemented for Android versions 2.2 and up, and will be evaluated in the European AAL-funded IS-Active project [3].

3.2 Content

The feedback message or task being communicated to the user is primarily categorised by its content. At a fundamental level there are two purposes for feedback, namely encouragement and discouragement of physical activity. A possible third type of feedback can be considered when the patient's activity is within accepted margins. Patients could eventually receive a third type of message as a status report to simply state that they are doing well; this type of message could prove motivating, but is currently left out of the Feedback Agent framework. The messages can then be further subdivided with respect to activity or task e.g. gardening, walking, sitting down, etc. These activities are either indoor or outdoor activities and encompass a range of instructions from leisurely activities to practical household tasks; ensuring that the available message ontology has a high number of different activities and allows for more variation in feedback as well as potential to cater closely for different individual lifestyles. Based upon answers to the preference questionnaires in the first phase of operation of the Feedback Agent (see Section 2), some patients may automatically become eligible for specific activity messages such as pet walking if they have a dog, or gardening if they have a garden, etc.

Once the content and general information to be communicated has been decided each activity message is then rewritten in multiple wordings, so as to maintain the information transferred but also introduce new added 'personality dimensions'. These dimensions are added style and tone qualities such as length of the message, argumentation quality, formality of tone, etc and are designed to transform the text from a basic message to a message personally tailored to

---

[3]http://www.is-active.eu/

match the needs of the recipient. The six chosen personality indicators were shown to be influential and significant based on a participant survey conducted. Based on the patient's previously answered questions and behaviour, in both the cold start phase and the real time phase, the system is able to qualitatively score the individual's openness to receiving instructions in a particular style, hereby amplifying the potential motivating effect and increasing the probability of a positive response. Using added personality dimensions allows for a higher volume of messages with the same base content, enabling more variety during selection.

Categorisation of the message based on content, context and style allows for a high volume message ontology wherein two users with the same behaviour but different personalities can be assessed as individuals. For example, one patient who enjoys clear instructions is told to "Go for a walk." whereas another is told "Moving is healthy, why don't you go for a walk?" due to his/her preference for more friendly and suggestive guidance. It is believed that boosting the personal relevance for each patient as a way to heighten motivating effect [1], will consequently also increase success of encouraged behaviour change.

In future, further information such as weather, GPS (with home/work location), local events, etc could be incorporated into the system for even more specific message content relevant to not only the individuals persona but also their lifestyle and surroundings.

The work described in Wieringa et al. describes the design of the software module for selecting the appropriate feedback message from the Message Ontology based on context-specific pruning of the ontology, and the use of a Compliance Database in which the user's preference to certain message content and styles are stored [13]. Parts of the message ontology that are irrelevant for the current user (e.g. **Outdoor Activities-Gardening** for users that do not own a garden) are removed from the Ontology, as well as certain activities that are inappropriate because of the user's current real-time context (e.g. **Outdoor Activities** in case of bad weather). Similar to the feedback timing module, responses to feedback messages are stored. The Compliance Database keeps track of how often the user responded positively and negatively to messages in each node of the Ontology, as well as how the user responded to the six different personality indicators. When choosing the message to present to the user, this information is then used to increase the probability of selecting messages to which the user responded well previously.

The work on message content selection is currently under development. The message ontology to use in evaluations is being constructed, and the software module is being adapted to work on the Android platform.

3.3 Representation

The third aspect of feedback on physical activity is the representation of the actual messages to the user. As can be seen in Figure 1, the current interface uses a simple textual representation of the messages. In order to provide the user with optimal usability and ensure an enjoyable long term human-machine interaction, better interface modalities and imple-

mentations should be explored. As the goal of our work is to create the Feedback Agent as a "personal activity coach", it seems obvious to look into human like representations of the interface. The field of Embodied Conversational Agents (ECAs) works on just that: human-like representations of software agents that can have conversations with a human being. The representation of an ECA can range in complexity from a simple picture (or avatar) to a fully animated 3 dimensional model of a human(like) being. The Mobile Fitness Companion developed by Ståhl et al. is an implementation of a mobile fitness application that can understand speech, and reply with spoken responses, thus implementing a conversational fitness coach [11]. The visual represenation on the Smartphone is a simple picture of a Nabaztag rabbit[4], mimicking the actual representation of the system's fixed home based system. A more advanced ECA has been implemented in the MOPET (Mobile Personal Trainer) application by Buttussi et al. [2]. Their virtual training companion, called Evita, depicted in Figure 3 is a fully animated 3D model that can guide users in doing fitness exercises by showing how to do them on the Smartphone screen; as well as motivate the users to maintain physical activity.

We are currently looking into the possibilities of implementing our Feedback Agent as an Embodied Conversational Agent, as we believe it can further improve the user's experience and increase the feeling of having a personal activity coach in your pocket.

Figure 3: Evita: a mobile 3D virtual fitness trainer aimed to guide and motivate users in doing fitness activities [2].

---

[4]http://www.nabzone.com/

4. DISCUSSION

We described our vision of an electronic personal activity coach that is context-aware, self-learning and personalized for individual users. The development of the smart generation of feedback timing and content is well underway and described in earlier work [9, 8, 13], while for the representation as Embodied Conversational Agent there is much research still to be done.

In the process of developing the Feedback Agent, certain issues have arisen regarding design and implementation, while more will possibly follow. We highlight here some of the issues already encountered and hope to receive valuable feedback from the community.

As mentioned in Section 3.2, the Feedback Agent currently works around the concepts of "encouraging" and "discouraging" feedback messages, that are given whenever the user needs to be *encouraged* to be more active or *discouraged* from it. The third type of "neutral" feedback, when the patient is doing well could be beneficial towards the overall goal of sustainable behaviour change. However, the learning aspects of the Feedback Agent is based on the notion of compliance to feedback. For *neutral* feedback messages, the notion of compliance is undefined, as the patient is not given a specific task. Implementing neutral messages, or compliments to patients who are performing according to their given reference is currently an open issue.

Another issue that we are currently unsure of how to handle lies with the notion of so-called "off-days". The Feedback Agent is constantly learning from the user's behaviour, but there are special cases that might better be ignored. When, for example, the patient is ill for a few days, he or she will probably suffer from a serious decline in physical activity as well as react differently to feedback. Similarly a user's behaviour can dramatically change simply because he or she is on holiday. In order to handle these cases efficiently and automatically, and not let them pollute the machine learners, an elegant solution will have to be found.

We hope to report our solutions and evaluations with the final system in the near future.

5. ACKNOWLEDGMENTS

6. REFERENCES

[1] F. C. Bull, M. W. Kreuter, and D. P. Scharff. Effects of tailored, personalized and general health messages on physical activity. *Patient education and counseling*, 36(2):181–92, Feb. 1999.

[2] F. Buttussi, L. Chittaro, and D. Nadalutti. Bringing mobile guides and fitness activities together: a solution based on an embodied virtual trainer. *Work*, pages 29–36, 2006.

[3] D. E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, 1989.

[4] J. Ho and S. S. Intille. Using context-aware computing to reduce the perceived burden of interruptions from mobile devices. *Proceedings of the SIGCHI conference on Human factors in computing systems CHI 05*, Portland,:909, 2005.

[5] M. W. Kreuter and V. J. Strecher. Do tailored behavior change messages enhance the effectiveness of health risk appraisal? Results from a randomized trial. *Health Education Research*, 11(1):97–105, 1996.

[6] A. Morandi and R. Serafin. A personalized motivation strategy for physical activity promotion in diabetic subjects. In *Proceedings of the 2nd Workshop on Personalisation for E-Health*, pages 51–55, Corfu, Greece, 2007.

[7] S. M. Noar, C. N. Benac, and M. S. Harris. Does tailoring matter? Meta-analytic review of tailored print health behavior change interventions. *Psychological Bulletin*, 133(4):673–693, 2007.

[8] H. op den Akker, H. J. Hermens, and V. M. Jones. A Context-Aware Adaptive Feedback System for Activity Monitoring. In *Proceedings of ICAMPAM 2011: the 2nd International Conference on Ambulatory Monitoring of Physical Activity and Movement*, 2011.

[9] H. op den Akker, V. M. Jones, and H. J. Hermens. Predicting Feedback Compliance in a Teletreatment Application. In *Proceedings of ISABEL 2010: the 3rd International Symposium on Applied Sciences in Biomedical and Communication Technologies*, Rome, Italy, 2010.

[10] J. O. Prochaska, C. C. DiClemente, and J. C. Norcross. In search of how people change: Applications to addictive behaviors. *American Psychologist*, 47(9):1102–1114, 1992.

[11] O. Stå hl, B. Gambäck, P. Hansen, M. Turunen, and J. Hakulinen. A Mobile Fitness Companion. *Science And Technology*, 2008.

[12] M. G. Van Weering. *Towards a new treatment for chronic low back pain patients*. PhD thesis, University of Twente, 2011.

[13] W. Wieringa, H. op den Akker, V. M. Jones, R. op den Akker, and H. J. Hermens. Ontology-Based Generation of Dynamic Feedback on Physical Activity. In *Proceedings of the 13th Conference on Artificial Intelligence in Medicine (AIME)*, pages 55–59, Bled, Slovenia, 2011. Springer.

[14] I. H. Witten and E. Frank. *Data Mining: Practical Machine Learning Tools and Techniques with Java Implementations*. Number San Francisco in The Morgan Kaufmann series in Data Management Systems. Morgan Kaufmann, 2000.