Ontology-Based Generation of Dynamic Feedback on Physical Activity

Wilko Wieringa[3], Harm op den Akker[1,2], Valerie M. Jones[2], Rieks op den Akker[3], and Hermie J. Hermens[1,2]

[1] Roessingh Research and Development
[2] University of Twente, Department of Telemedicine
[3] University of Twente, Department of Human Media Interaction

Abstract. Improving physical activity patterns is an important focus in the treatment of chronic illnesses. We describe a system to monitor activity and provide feedback to help patients reach a healthy daily pattern. The system has shown positive effects in trials on patient groups including COPD and obese patients. We describe the design and implementation of a new feedback generation module which improves interaction with the patient by providing personalised dynamic context-aware feedback. The system uses an ontology of messages to find appropriate feedback using context information to prune irrelevant paths. The system adapts using derived probabilities about user preferences for certain message types. We aim to improve patient compliance and user experience.

1 Introduction

Balancing daily activity is a focus in the treatment of chronic diseases such as Chronic Low Back Pain (CLBP) and Chronic Fatigue Syndrome (CFS) [1], and COPD and obesity. Keeping an active lifestyle, while keeping in mind the limits caused by the disease and the risks of exacerbations, is a major goal in preventing COPD patients from entering a downward spiral of inactivity and disease progression. A popular approach for promoting an active lifestyle is to objectively measure the subject's activity pattern and provide feedback periodically, or at the end of each day. We describe an activity monitoring system which has been trialled on different patient groups including COPD and obese patients. Experience with this system, and reported elsewhere in the literature [2, 3] highlights the need for a more sophisticated approach to the delivery of user feedback. Not only the content of the message, but the timing and manner in which it is delivered need to be tuned to the individual and their context. The expression of message content too needs to be varied in order to avoid habituation effects, automatic responses, boredom or irritation. The activity monitoring system consists of a 3D-accelerometer connected via Bluetooth to a PDA. The system logs acceleration every 10 seconds as an integrated value, summed over the three axis of movement. The patient's current activity level is compared to a predefined reference, and feedback is given at fixed hourly intervals: "*Encouraging*" to encourage activity, "*Discouraging*" to discourage activity and "*Neutral*" feedback if the patient is doing well.

2 Improving Feedback Message Generation

The objective of the work described here is threefold. First, we aim to bring more variation to the feedback messages presented. Hearing the same messages over and over again reduces the subject's willingness to use the system and comply to the advice given in the feedback [4]. Second, we want to generate messages that are relevant to the subject's current context or environment, e.g. when the weather is good, advise people to go for a walk, when the weather is bad, advise an indoor activity. Third, we want to create a system that can be tuned to the user's personal preferences. If one individual responds better to messages that are given in a commanding tone, the system should adapt to this. In short, the system should be *dynamic*, *context-aware* and *self adapting*.

2.1 Dynamic Message Generation

At the core of the feedback message generation system lies an ontology of feedback messages. Erriquez and Grasso [2] show the viability of this approach over more complex natural language generation methods. Key factors in their approach are scalability and opacity, as for any generated message the reason why that message was generated from the ontology can be deduced. Our OWL2.0[4] ontology contains a set of feedback messages structured in a meaningful way. When the system requests a feedback message, the ontology is traversed, generating a list of possible candidate sentences from which a random selection is made. Initially we generate three message types: "*Neutral*" (when the subject is performing well), "*Encouraging*" (meaning encouraging more activity when the subject is insufficiently active) and "*Discouraging*" (meaning discouraging activity when the subject is too active). We used a threshold of 10% compared to the reference to determine over- or under performing. Each of these *MessageType*s can be composed of several *MessageComponent*s, the top-level entity for all the actual sentences in the ontology. A MessageType currently consists of one *Evaluation* (a message containing factual information on the current activity performance) and either one of *DiscouragingAdvice*, *NeutralAdvice* or *EncouragingAdvice*. An example of a generated EncouragingMessage could be "You have been insufficiently active. [Evaluation] Please go for a walk! [EncouragingAdvice]". We focus here on the advice components as they contain the motivational messages that are of most interest to the physical activity monitoring system. Figure 1 shows part of the message ontology for *EncouragingAdvice*, with the *EncouragingInside* and *Household* entities expanded. The lowest level contains the actual feedback messages (grey circles). At runtime the system traverses this ontology, gathering all the candidate sentences and makes a random

[4] http://www.w3.org/TR/owl2-overview/

choice from the message candidates for presentation to the user. This initial version of the system is equivalent to selecting a random message from a list of predefined feedback messages; however, it forms the basis for the next version of our context-aware and self-adapting system, explained in the two following sections.

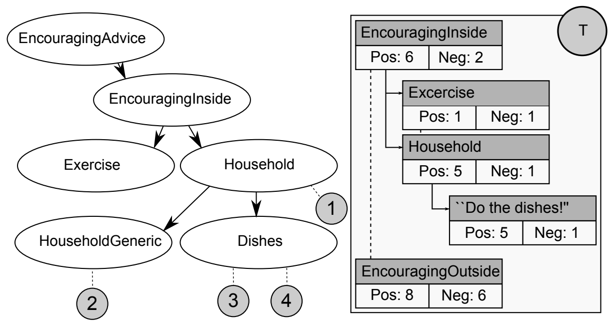

Fig. 1. (Left) Example of the ontology with EncouragingAdvice expanded and individual feedback messages at the leaves. (Right) Compliance database for storing responses (compliance) to feedback messages (see Section 2.3).

2.2 Adding Context-Awareness

Notice in Figure 1 there is a high level split between messages that tell the patient to go outside (*EncouragingOutside*), and those that tell the patient to stay inside (*EncouragingInside*). However a patient should not be sent outside if in a thunderstorm; hence our system is extended to include Boolean functions attached to ontology entities. In this case the *EncouragingOutside* entity has the function `isWeatherGood()` attached, which returns true if the weather at the patient's location can be considered "good". In our Android implementation, this function is implemented by retrieving current location and weather data online and performing a rule-based conversion to Boolean. Similar context functions, using for example location information, could be added to rule out a branch of messages that are only relevant when the patient is at home, or at work. The same schema can also be used to embed user preferences into the selection process for message candidates. If user information, such as dog ownership, is available this can be used to open up branches containing messages related to taking the dog for a walk.

2.3 A Self-Adapting Mechanism for Message Selection

Instead of choosing randomly between messages candidates deemed to be relevant given the message ontology and the current contextual state, we can learn from the patient's response to previously given feedback messages. We use the same definition of feedback compliance as [5]; the compliance score "*compares the amount of activity performed in the 30 minute interval before the feedback event (Δ_1), with the amount of activity performed in the 30 minute interval after the feedback event (Δ_2)*". The 30 minute interval was chosen for comparability with earlier work, but this value can be seen as a variable that needs further study. For encouraging messages, the subject complied if $\Delta_1 < \Delta_2$, and for discouraging message the subject complied if $\Delta_1 > \Delta_2$. Thirty minutes after each feedback message, the compliance score (`true` or `false`) is fed back to the system. Notice that we look at the time the subject has seen the feedback message for calculating compliance. The number of times each individual message was judged positively and negatively is stored in a `compliance database`. These compliance scores are used to make an informed decision about which message to select. The winner selection algorithm is explained, using the choice between *EncouragingInside* and *EncouragingOutside* in Figure 1 (right) at some state T as an example. We first define $levelData$ as the total amount of data available on the *level* in the ontology of our choice. In this case $levelData = 6+2+8+6 = 22$. Furthermore, we define $localTrue$ as the number of positive judgments for each individual entity, and $localData$ as its total amount of judgements. Then:

$$entityChance(X) = \frac{localTrue}{localData} \times (1 - \frac{localData}{levelData}) \tag{1}$$

$$entityChance(EncouragingInside) = \frac{6}{8} \times (1 - \frac{8}{22}) = \frac{21}{44} = \frac{147}{308} \tag{2}$$

$$entityChance(EncouragingOutside) = \frac{8}{14} \times (1 - \frac{14}{22}) = \frac{16}{77} = \frac{64}{308} \tag{3}$$

After computing the common denominator for each fraction, we define the weight of each entity as their fraction numerator, and their selection probability as $\frac{entityNumerator}{sumOfNumerators}$ (e.g. the selection probability for EncouragingInside is $\frac{147}{211} \approx 70\%$ and for EncouragingOutside $\frac{64}{211} \approx 30\%$). Next we define a lottery, where a total of $100 \times N$ (where N is the number of entities on the current level) balls are divided among the candidates. First a fixed percentage of 20% of the balls is equally divided amongst all candidates, the rest of the balls are divided according to the weights calculated above. In this case: $EncouragingInside = 20 + 160 \times \frac{147}{211} = 131$ and $EncouragingOutside = 20 + 160 \times \frac{64}{211} = 69$. The algorithm has the following three properties: (1) it selects with a higher probability those candidates which have been judged positively more often than negatively (2) it selects with a higher probability those candidates for which less data is available, and (3) there always remains a small probability that a candidate is selected, regardless of positive/negative ratio. These properties ensure that the algorithm (1) learns to give (types of) feedback messages that the individual prefers more often, (2) gathers enough data on all messages to make informed decisions, and (3) does not fully discard messages that so far have a 100% negative judgment.

2.4 Evaluation

In order to evaluate the properties of the system we ran simulations using fictional personas. We generated 300 feedback messages[5] and determined the patient's compliance to those messages on three personas: "sporty", "random" and "housefather". The sporty persona had a high probability of reacting well to sporting activities, while the "housefather" complied better to household activities. After 300 messages, only 18% of the messages generated for the non-random personas were of a disliked category. This shows that adaptation is taking place, although evaluations with real patients is a crucial next step.

3 Discussion

We have described the design and preliminary testing of a feedback message generation component of an activity monitoring system that has been integrated into our Android based framework. Simulation runs indicate that the self-adapting mechanisms work as intended. The next step is to test the system on real patients. Current work focusses on designing an ontology that can be used in large patient trials; the content of the ontology will be defined based on the literature on activity based feedback and on results of patient questionnaires. The ontology will be structured in a logically and semantically sound way in order to optimize learning and context-awareness. By combining the system described here with our previous work on designing a self-learning algorithm for determining the optimal timing for feedback messages for individual patients [5], these two components of the activity monitoring system form the cornerstone for an intelligent Feedback Coach. The coach will be able to act as a virtual companion for patients who try to change their behaviour and lead a more healthy lifestyle. The generic design approach and focus on personalisation also means that this application is not *limited* to patients suffering from chronic diseases.

References

1. Van Weering, M.G., Vollenbroek-Hutten, M.M., Tönis, T.M., Hermens, H.J.: Daily physical activities in chronic lower back pain patients assessed with accelerometry. European journal of pain London England **13** (2009) 649–654
2. Erriquez, E., Grasso, F.: Generation of Personalised Advisory Messages : an Ontology Based Approach. In: Proceedings of the 21st IEEE International Symposium on Computer-Based Medical Systems (CMBS), Jyvaskyla, Finland (2008) 437–442
3. Cortellese, F., Nalin, M., Morandi, A., Sanna, A., Grasso, F.: Personality Diagnosis for Personalized eHealth Services. In: Proceedings of the 2nd International ICST Conference on Electronic Healthcare for the 21st Century, Springer (2009) 173–180
4. Bickmore, T.W., Picard, R.W.: Establishing and maintaining long-term human-computer relationships. ACM Transactions on CHI **12** (2005) 293–327
5. op den Akker, H., Jones, V.M., Hermens, H.J.: Predicting Feedback Compliance in a Teletreatment Application. In: Proceedings of ISABEL 2010, Rome, Italy (2010)

[5] 300 is approximately the number of messages that patients received in one month's time in earlier studies.