Maintaining Engagement in Shared Goals with a Personal Robot Companion through Motivational State Modeling

Elena Corina Grigore
Department of Computer Science
Yale University
New Haven, CT 06520, USA
elena.corina.grigore@yale.edu

Brian Scassellati
Department of Computer Science
Yale University
New Haven, CT 06520, USA
brian.scassellati@yale.edu

ABSTRACT

This work focuses on developing adaptive human-robot interaction systems that work with users over long periods of time to achieve a common goal that is beneficial to the user. The system adapts to the user by modeling his or her motivational state in order to achieve better cooperation during the interaction, which in turn could lead to higher levels of goal achievement. The particular scenario we focus on is that of a robot companion interacting with adolescents, helping them succeed at achieving daily physical activity goals. The robot does so by modeling the user and adapting motivational strategies best suited for each user. The key question we are investigating within this collaboration context is: how can the robot use long-term interaction to create and maintain a user model of its interaction partners that allows for such an adaptation? These models allow for the interpretation of "why" and "how" a user succeeded or failed at achieving a physical activity goal, and are based on both physical activity data obtained from wearable sensors (such as wristband devices) and information acquired by the robot from its interaction partners. This interpretation is central to an adaptive human-robot interaction system that socially manipulates users towards shared goals.

1. INTRODUCTION

Cooperation is central to human relationships and societal structures and is based on social norms that can be applied in a peer group, family, organization, or entire society [11, 22]. Cooperative behaviors are at the heart of human existence, and are beneficial to individuals and to society as a whole. Most cooperative behavior tasks assume both parties are equally motivated and involved in achieving a common goal. Some cooperative behavior tasks, however, face the challenge of having one of the parties become less motivated over time even when having the same high level goal. Examples include people who work together with a weight loss coach with the purpose of loosing weight, or people who work together with a personal trainer to reach a fitness goal. Even though they engage in cooperative behavior in order to reach a goal that personally benefits them, people become less motivated during the process of achieving it. Psychology research shows that cooperative behavior can be predicted and that positive emotion and low levels of inhibition are important for achieving higher levels of cooperation [22]. Modeling users' motivational state is thus of paramount importance if we are to develop adaptive robots that aim to help users achieve goals through cooperative interaction. In our current work, we focus on modeling users' motivational state in the context of helping them stay engaged in high levels of physical activity.

The benefits of physical activity are well known, as research shows that daily physical activity has wide-ranging benefits, from improving cognitive and academic performance [28] to helping with bone development and health [7]. Our work seeks to sustain these benefits with robot home companions through personalized, data-driven coaching.

The first comprehensive guidelines on physical activity for individuals ages 6 and older was released in 2008 by the U. S. Department of Health and Human Services. The guidelines state that children and adolescents should aim to accumulate at least 60 minutes of moderate- or vigorous-intensity aerobic physical activity on most days of the week, preferably daily [1]. Current evidence shows that levels of physical activity among youth remain low, and that levels of physical activity decline dramatically during adolescence [2]. In 2008 it was found that only 8% of adolescents were active in moderate- to vigorous-intensity activity on 5 days per week for at least 60 minutes each day [27]. These data show the importance of developing methods to keep adolescents on track to achieving daily recommended levels of physical activity. With this in mind, we seek to develop a human-robot interaction system that motivates adolescents to engage in physical activity on a daily basis.

2. BACKGROUND AND RELATED WORK

There exists a great deal of work concerning health and well-being applications, such as commercial systems that support goal-setting and reflection for the user, and employ motivational strategies for achieving goals [18], [12]. Other prior work explores tracking of long term goals along with factors that have an important role in goal-setting theory [5].

In the field of HCI, wearable sensors are being used in the development of persuasive technologies that motivate healthy behavior. Work in this area is wide, ranging from activity recognition systems that employ this information to keep users engaged in physical activity [8] to how such applications should be evaluated [14]. Work related to contextual information playing a role in daily physical activity is scarce, and focuses on showing that providing users with such information (activity, location, people) increases their awareness of physical activity [17].

Social robots have also been used to promote and keep users engaged in physical activity. Work in this area touches on investigating the role of praise and relational discourse and that of physical vs. virtual embodiment [10], on how maintain engagement in the weight loss domain [13], and on how users judge the robot's capabilities and competence when it takes on the role of a fitness instructor vs. that of a social co-participant [25]. The feasibility of using a SAR (Socially Assistive Robotics) approach to keep children aged 5 - 8 engaged in a three-week long interaction with a robot that taught them how to make healthy food choices has also been investigated [23]. The study highlights that children engaged with the robot throughout the interaction, responding to its questions and becoming engaged with the task presented to them.

To date, however, there is no system that ties together continuous remote monitoring of user physical activity obtained from the wearable sensors with an assistive robot for daily, long-term interactions. This link consists of the user model that interprets the state of the robot's interaction partner in order to keep him or her on track to achieving the shared goal. The interpretation obtained from the user model stands as a basis for the adaptive system to learn which strategies work best for an individual user.

The user model we are building employs an ontology-based approach. Ontologies are defined as explicit accounts of shared understanding in a given subject area and bolster communication between users, experts, and technology systems [29]. Since ontologies are extremely useful in providing a formalism that specifies and clarifies concepts in a domain and the relationships among them, their value for health applications has been recognized by different lines of research. Such research includes the investigation of computational models of dialogue trying to simulate a human health counselor [6] and that of computerized behavioral protocols (CBPs) that help individuals improve their behaviors [16]. The former focuses on automating dialogue in order to best simulate a human counselor, and is evaluated via questionnaires with respect to how close it comes to emulating a counselor's empathy, naturalness, trust, etc. The latter focuses on creating an ontology useful for modeling PACE-Adolescent, a behavioral protocol aimed at promoting healthy physical activity and dietary behaviors in adolescents. To date, there does not exist a user model based on an ontology which helps describe and interpret "why" and "how" the user succeeded or failed at a given physical daily goal.

3. METHODOLOGY AND DESIGN OF THE SYSTEM

The robot platform chosen for this study is Keepon, a non-mobile platform with four degrees of freedom, designed to interact with children [15]. We are using a version of the My

Figure 1: Keepon Robot

Keepon toy, modified to be programmable, that can be seen in Figure 1, above. The adolescent wears a wristband device that keeps track of the number of steps taken throughout the day. He or she interacts with a robot once daily, both directly and via a phone application. The application presents an avatar for the robot - a virtual character with similar appearance to the robot - with simple, 2D animations. The robot communicates with the user via pre-scripted snippets of speech, using a text-to-speech module. The participant is then given a choice between pressing a microphone icon to speak to the robot or inputting an answer to the robot's question using the application interface (e.g. if the user is being asked for their name, they can press the microphone icon and say their name or they can press a button which brings up a keyboard for input). If the system cannot recognize the speech, the application defaults to its native interface for entering information.

The robot has a back-story, which unfolds over time in order to keep adolescents engaged throughout the interaction. The story is that the robot is a robot-alien, named EphyT, that landed on Earth and needs the adolescent's help to return home. The closer the user gets to accomplishing daily physical activity goals, the more energy EphyT gains. Thus, the robot and the adolescent engage in a collaborative interaction, working together towards both helping the user engage in high levels of physical activity and helping the robot get back to its home planet.

The system is depicted in Figure 2, on the next page. When the adolescent interacts with the robot, the back-story unfolds. The robot then asks the user a series of questions. These questions are meant to obtain extra information (described below in the User Model section) about the user in order to build and maintain a user model. The system then acquires extra information from online sources about external factors (e.g. weather or school schedule announcements) and feeds both this and the user's answers to questions as inputs to the user model. The output of the user model, which is discussed below, then becomes part of the input to the adaptive system. The other input to the adaptive system is the physical activity data the robot acquires from the wearable sensor. This information is all fed into the adaptive system, described in the Motivational Strategies section below, which outputs an appropriate motivational strategy for the user. This strategy is used to shape the new physical activity goal the adolescent needs to accomplish the next day.

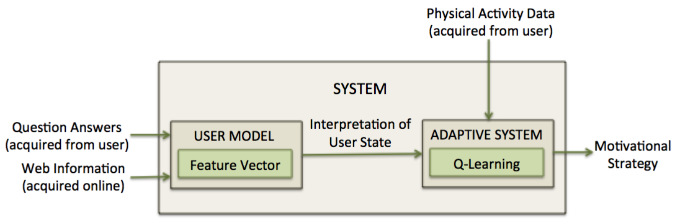

Figure 2: System Diagram

4. USER MODEL

The ontology our model relies upon needs to represent the core concepts employed by professional trainers and health counselors related to creating a profile of the user and keeping track of how this changes over time. The creation and validation of ontologies relies heavily on expert interviews and validation. In this work, a user model is created based on key elements influencing the achievement of a goal. These are factors long studied in goal setting theory [20], social cognitive theory [4], self determination theory [9], and theory of planned behavior [3], and are discussed in the paragraph below. The validation and identification of missing core concepts and the relationships that exist among them is realized via initial expert interviews, including personal fitness trainers and health counselors. The user model is used by the human-robot interaction system in order to adapt to a particular user and choose the appropriate motivational strategy for a user at every time step. A time step is defined as a day since the robot interacts with the adolescent once daily and sets daily physical activity goals.

The user model is based on concepts identified from literature review, as previously discussed. As can be seen in Figure 2, the inputs for the user model are dichotomized into two categories. The first represents the user's answers to questions asked by the robot about how the participant felt while trying to accomplish the goal for the day. These questions ask the participant about the key elements influencing the achievement of a goal, namely socio-structural factors, social pressure, self-efficacy, and attitude toward behavior. They are phrased in an easy to understand manner, avoiding academic and arcane terms. *Socio-structural factors* are external elements that might have contributed to the success or failure at a given goal and are classified into facilitators and impediments, e.g. an adolescent's particular schedule for the day or items from the schedule that might act as either a facilitator or an impediment in accomplishing the physical activity goal. *Social pressure* models how much pressure the adolescent felt while trying to accomplish the goal for the day, given the particular motivational strategy last employed by the robot. *Self-efficacy* models how confident the adolescent felt while trying to accomplish the goal. Finally, the *attitude toward the behavior* models how much commitment the user had toward accomplishing that particular goal.

The second input category represents external information acquired by the system online about factors that might have influenced the achievement of a goal. Such factors are information about the weather for the day or announcements about school schedule changes and cancellations that the system can acquire directly, without the need to ask the user. This information is publicly available and can be acquired, for example, from the school's website or public weather reports. This information is fed into the socio-structural factors, into either the facilitators or impediments feature. All of these core factors make up the feature vector for a user, which contains numerical values based on the user's responses during the interaction with the robot.

The output of the user model is an interpretation of the user's state. As will be discussed in the next section, this interpretation represents the state of the world (the state the user is in) at that particular time step. The numerical value of each feature is mapped onto one of three discrete categories representing intensity levels, namely low, neutral, and strong. This mapping is applied to reduce the state space size and produce a more intuitive interpretation of the user's state. For example, a numerical value of 7 on a 1-to-7 scale for self-efficacy would be assigned to the "strong" intensity level. While a finer grained representation may provide a more accurate depiction of the state of the world, collecting sufficient data to exploit this specificity isn't feasible in the current work. This would impose unrealistic sampling constraints on our approach given that we obtain a limited number of samples from participants and that the action set we are using (the number of motivational strategies) is quite small. The interpretation of a user's state thus consists of a feature vector containing intensity levels for each feature. This interpretation is of paramount importance for an adaptive human-robot interaction system whose aim is to keep the user engaged in accomplishing a shared goal each day since it represents the basis for choosing an appropriate motivational strategy for the user.

In order to validate this user model and create an ontology which links the core elements modeling a user to motivational strategies employed by experts, we utilize a formal interview process. The interview thus aims to validate and identify missing main concepts in the domain of motivation for physical exercise and the relationships among them. The interview is structured in two parts, as follows. The first part asks the expert to list (1) important factors for keeping students engaged in physical activity (corresponding to the main concepts used in the ontology), (2) techniques used by experts to identify the students' motivational state (used to create a mapping from users' answers to questions to the state-space), (3) information about students that helps the expert interpret the students' success or failure (used to validate the user model's output), and (4) strategies and techniques employed by the personal trainer to keep students engaged in physical activity (corresponding to the motivational strategies discussed above). The second part presents the expert with the same categories, but with given answers (drawn from literature). The expert is asked to indicate how much he or she agrees with the specific answer, using a Likert scale. The goal of this initial interview is to check if experts indeed use concepts identified in literature as key factors, acquire new key factors used in practice, and acquire the basis for creating an ontology linking the user model to the motivational strategies used by experts.

5. MOTIVATIONAL STRATEGIES

Our adaptive human-robot interaction system bases its decisions of what motivational strategy to employ when setting a daily physical activity goal on the user model's interpretation of "why" and "how" the user was (un)successful at his or her previous goal, i.e. the interpretation of the user's state. This is possible since the motivational strategies identified from literature are associated with factors present in the user model, e.g. "Intrinsic motivation is associated with the desire to master a task and the satisfaction that comes with that, whereas extrinsic motivation refers to completing tasks solely as an ego boost, be it beating peers in a competition, or receiving praise from a parent, teacher or colleague" [21].

Two main motivational strategies identified from literature are strategies emphasizing cooperation and competition. Cooperation strategies include setting out a physical activity goal in a way that fosters cooperation between the user and the robot. In [24] cooperative exercise game players lost significantly more weight than players in the control condition, who gained weight over time. Competition can and has been shown to be effective within exercise game interventions, e.g. for most participants in [19], competitiveness presented a more stimulating challenge than cooperation. Some of the participants, however, felt that competitiveness was incompatible with the spirit of the game, creating the need to develop an adaptive system that can model the preferences of each user. Other motivational strategies are currently being researched and include making the user aware of the importance of engaging in physical activity and of the negative health consequences of not exercising, autonomy support, structure, and involvement [26].

Our system adapts to an individual user by following a Q-learning approach [30]. Q-learning can be used when an agent wants to learn an optimal policy from its history of interaction with the environment. The agent can typically choose from a finite collection of actions at every time step. Q value functions are state-action pair functions that estimate how good a particular action will be in a given state, i.e. what the return for that action is expected to be. The value of taking action a in state s under a policy π is noted as $Q^{\pi}(s, a)$ and represents the expected return when starting from state s, taking action a, and thereafter following policy π.

In the current work, the actions the "agent" (in our case the robot) can take are the different motivational strategies discussed above, $a \in \{m_1, m_2, m_3, m_4, m_5, m_6\}$. The states of the world are the interpretations of a user's state, as explained in the User Model section. Thus, a state is a feature vector representing the interpretation of the user's state, and is defined as $intrp_t = [f_1\ f_2\ f_3\ f_4\ f_5]$, where the f_is represent the features discussed above, which take values associated with intensity levels. The reward for a particular state is the difference between the number of steps taken by the user and the number of steps set as the goal, $r_t = \#steps_{taken} - \#steps_{goal}$. At time step t, the algorithm observes the current state $intrp_t$, chooses an action a_t among the motivational strategies based on an ϵ-greedy policy, takes the action, observes the reward r_t as well as the new state $intrp_{t+1}$, and then updates the Q-value for the state using the observed reward and the maximum reward possible for the next state. The update is performed based on the following formula: $Q(intrp_t, a_t) = Q(intrp_t, a_t) + \alpha[r_t + \gamma \max_{a_{t+1}} Q(intrp_{t+1}, a_{t+1}) - Q(intrp_t, a_t)]$, where α and γ are both set between 0 and 1 and specify how quickly learning occurs and how much future rewards matter, respectively. The algorithm will thus work toward finding an optimal policy, in order to maximize the expected return, i.e. positive reward or small values for negative rewards.

6. FUTURE WORK

Future work involves validating the ontology and our ontology-based user model. The system itself is to be validated as part of a user study. The study will employ the adaptive human-robot interaction system vs. a non-adaptive system that interacts with adolescents in the same manner, daily, but without adapting to the user.

7. REFERENCES

[1] 2008 physical activity guidelines for americans. Technical report, U. S. Department of Health and Human Services, 2008.

[2] Physical activity guidelines for americans midcourse report: Strategies to increase physical activity among youth. Technical report, U. S. Department of Health and Human Services, 2013.

[3] I. Ajzen. The theory of planned behavior. *Organizational behavior and human decision processes*, 50(2):179–211, 1991.

[4] A. Bandura. Social cognitive theory of self-regulation. *Organizational behavior and human decision processes*, 50(2):248–287, 1991.

[5] D. Barua, J. Kay, B. Kummerfeld, and C. Paris. Modelling long term goals. In *User Modeling, Adaptation, and Personalization*, pages 1–12. Springer, 2014.

[6] T. W. Bickmore, D. Schulman, and C. L. Sidner. A reusable framework for health counseling dialogue systems based on a behavioral medicine ontology. *Journal of biomedical informatics*, 44(2):183–197, 2011.

[7] C. Boreham and C. Riddoch. The physical activity, fitness and health of children. *Journal of sports sciences*, 19(12):915–929, 2001.

[8] S. Consolvo, D. W. McDonald, T. Toscos, M. Y. Chen, J. Froehlich, B. Harrison, P. Klasnja, A. LaMarca, L. LeGrand, R. Libby, et al. Activity sensing in the wild: a field trial of ubifit garden. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1797–1806. ACM, 2008.

[9] E. L. Deci and R. M. Ryan. Self-determination theory: A macrotheory of human motivation, development, and health. *Canadian Psychology/Psychologie canadienne*, 49(3):182, 2008.

[10] J. Fasola and M. J. Matarić. Socially assistive robot exercise coach: Motivating older adults to engage in physical exercise. In *Experimental Robotics*, pages 463–479. Springer, 2013.

[11] E. Fehr and U. Fischbacher. Social norms and human cooperation. *Trends in cognitive sciences*, 8(4):185–190, 2004.

[12] HealthMonth. Healthmonth - live healthier, for fun!, 2014. Available at `https://www.healthmonth.com/`.

[13] C. D. Kidd. Designing for long-term human-robot interaction and application to weight loss. 2008.

[14] P. Klasnja, S. Consolvo, and W. Pratt. How to evaluate technologies for health behavior change in hci research. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 3063–3072. ACM, 2011.

[15] H. Kozima, M. P. Michalowski, and C. Nakagawa. Keepon. *International Journal of Social Robotics*, 1(1):3–18, 2009.

[16] L. Lenert, G. J. Norman, M. Mailhot, and K. Patrick. A framework for modeling health behavior protocols and their linkage to behavioral theory. *Journal of biomedical informatics*, 38(4):270–280, 2005.

[17] I. Li, A. K. Dey, and J. Forlizzi. Using context to reveal factors that affect physical activity. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 19(1):7, 2012.

[18] Lifetick. Lifetick - goal setting the way it should be, 2014. Available at `https://www.lifetick.com/`.

[19] J. J. Lin, L. Mamykina, S. Lindtner, G. Delajoux, and H. B. Strub. Fish'n'steps: Encouraging physical activity with an interactive computer game. In *UbiComp 2006: Ubiquitous Computing*, pages 261–278. Springer, 2006.

[20] E. A. Locke and G. P. Latham. Building a practically useful theory of goal setting and task motivation: A 35-year odyssey. *American psychologist*, 57(9):705, 2002.

[21] J. G. Nicholls. Achievement motivation: Conceptions of ability, subjective experience, task choice, and performance. *Psychological review*, 91(3):328, 1984.

[22] D. G. Rand, G. Kraft-Todd, and J. Gruber. The collective benefits of feeling good and letting go: Positive emotion and (dis) inhibition interact to predict cooperative behavior. *PloS one*, 10(1):e0117426, 2015.

[23] E. Short, K. Swift-Spong, J. Greczek, A. Ramachandran, A. Litoiu, E. C. Grigore, D. Feil-Seifer, S. Shuster, J. J. Lee, S. Huang, et al. How to train your dragonbot: Socially assistive robots for teaching children about nutrition through play. In *RO-MAN, 2014 IEEE*. IEEE, 2014.

[24] A. E. Staiano, A. A. Abraham, and S. L. Calvert. Adolescent exergame play for weight loss and psychosocial improvement: a controlled physical activity intervention. *Obesity*, 21(3):598–601, 2013.

[25] L. Sussenbach, K. Pitsch, I. Berger, N. Riether, and F. Kummert. Can you answer questions, flobi?: Interactionally defining a robot's competence as a fitness instructor. In *RO-MAN, 2012 IEEE*, pages 1121–1128. IEEE, 2012.

[26] I. M. Taylor and N. Ntoumanis. Teacher motivational strategies and student self-determination in physical education. *Journal of Educational Psychology*, 99(4):747, 2007.

[27] R. P. Troiano, D. Berrigan, K. W. Dodd, L. C. Masse, T. Tilert, M. McDowell, et al. Physical activity in the united states measured by accelerometer. *Medicine and science in sports and exercise*, 40(1):181, 2008.

[28] F. Trudeau and R. J. Shephard. Physical education, school physical activity, school sports and academic performance. *International Journal of Behavioral Nutrition and Physical Activity*, 5(1):10, 2008.

[29] M. Uschold and M. Gruninger. Ontologies: Principles, methods and applications. *The knowledge engineering review*, 11(02):93–136, 1996.

[30] C. J. Watkins and P. Dayan. Q-learning. *Machine learning*, 8(3-4):279–292, 1992.