

On Designing a Social Coach to Promote Regular Aerobic Exercise

Shiwali Mohan, Anusha Venkatakrishnan, Michael Silva, Peter Pirolli

Interactive Intelligence Area, Palo Alto Research Center
Palo Alto, California 94306

{shiwali.mohan, anusha.venkatakrishnan, michael.silva, peter.pirolli}@parc.com

Abstract

Our research aims at developing interactive, social agents that can coach people to learn new tasks, skills, and habits. In this paper, we focus on coaching sedentary, overweight individuals to exercise regularly. We employ adaptive goal setting in which the coach generates, tracks, and revises personalized exercise goals for a trainee. The goals become incrementally more difficult as the trainee progresses through the training program. Our approach is model-based - the coach maintains a parameterized model of the trainee's aerobic capability that drives its expectation of the trainee's performance. The model is continually revised based on interactions with the trainee. The coach is embodied in a smartphone application which serves as a medium for coach-trainee interaction. We show that our approach can adapt the trainee program not only to several trainees with different capabilities but also to how a trainee's capability improves as they begin to exercise more. Experts rate the goals selected by the coach better than other plausible goals, demonstrating that our approach is effective.

Introduction

Intelligent social agents have tremendous potential to make human lives better by supporting us in our daily tasks, managing our daily schedules, and helping us learn new, complex skills. This paper focuses on the design and analysis of a particular type of social agents - a coach. The primary role of a coaching agent in a human-agent collaborative setting is to help the human trainee gain knowledge, skills, and tools to perform a new task. The coach may also motivate a trainee to strive for challenging variations of the task and/or provide emotional support in case of continued failures.

If intelligent agents are to be successful in coaching a person, they must take into account the person's specific needs, circumstances, and capability in their reasoning and decision making. People vary greatly along these factors and these factors evolve over time with experience with tasks. It is critical, then, that a coaching agent must represent a person's state describing these factors as well as how it changes over time. The coaching agent must tailor its coaching strategy to each specific person (*personal adaptation*) as well as to how a person evolves while training with a coach (*temporal adaptation*). Training a person for a new task may take

a long time (2-3 months) and several sessions. A coaching agent, therefore, is required to be a long-living system that maintains an ongoing interaction with its trainee.

The coach described here is designed to train overweight, sedentary individuals to develop capability and strength for regular aerobic exercise. Engaging in regular aerobic exercise such as biking, walking etc. increases overall energy expenditure above and beyond resting energy expenditure. This helps maintain a healthy weight as well as improve outcomes for weight-related co-morbidities such as type II Diabetes Mellitus, dyslipidemia etc. (Garber et al. 2011) that affect a large number of people around the world. To alleviate these problems, physicians recommend walking more. It is simple, versatile, requires limited resources, and is easily adaptable to individuals with varying capabilities. However, people need support for selecting specific goals to work on and close monitoring to develop exercising habits. A coaching agent for walking will make this support available for a large number of people at affordable cost.

This paper focuses on the goal setting strategy of coaching that is employed in human-human training scenarios (Shilts, Horowitz, and Townsend 2004). Here, the coach sets relevant and appropriate goals for their trainee. To be effective, a coach must set goals that are *difficult yet attainable*. Goals must induce effort for a trainee to improve performance but should not be too difficult to be successful at. If a trainee walks 15 minutes every day, a difficult yet attainable goal could be to walk 25 minutes every day. Additionally, the goals must be specific providing a clear and narrow target for which the required amount of effort can be estimated. Goals must be proximal mobilizing effort in the near future. Long-term, distal goals make it easy to postpone effort. *Eat two vegetables at lunch tomorrow* is more effective than *eat more vegetables* which is both general and distal.

The coach is embodied in a smartphone application through which it interacts with its human trainee. In interactions, it measures the trainee's current state, recommends exercise goals, and evaluates the trainee's performance. To set effective goals, the coach maintains hypotheses about the trainee's current aerobic capability. It employs a *trainee model* which is continually revised based on how the trainee performs on recommended goals. The coach's recommendations are then heuristically biased by the model's estimation.

The paper contributes the following. We propose an ab-

Activity	Detail	<i>i</i> (MET)
Moderate walking	1 mile in 30 – 45 minutes	3.3
Brisk walking	1 mile in 15 – 20 minutes	6
Interval brisk walking	5 minute moderate + 1 minute brisk walk (1 mile in 20 – 35 minutes)	4.5
Duration (<i>d</i> minutes)		
<i>max</i>	<i>min</i>	<i>increment</i>
30	10	5
Frequency (<i>f</i> days/week)		
<i>max</i>	<i>min</i>	<i>increment</i>
5	3	1

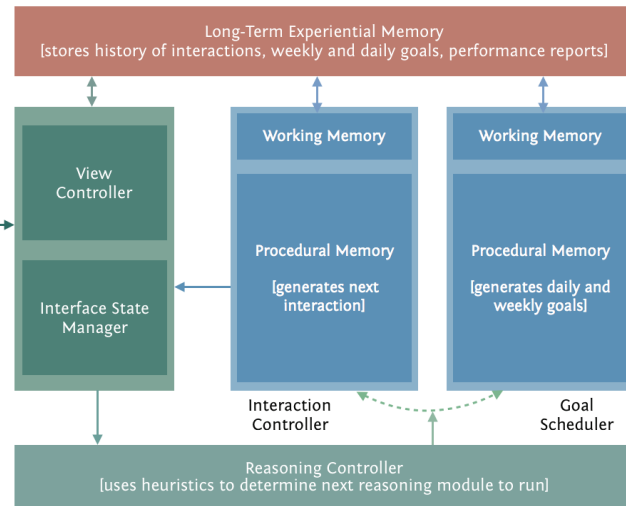
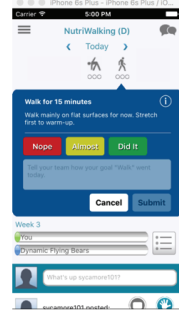


Figure 1: (Left) Goals considered by the coach: various walks and their intensity (*i*), valid values for duration (*d*) of a session, and valid values for number of sessions (*f*) in a week. (Right) Interactive interface and architecture of the social coach.

struct, parameterized model for growth in aerobic capability which encodes factors that experts use in their prescription. It affords online revisions by changing parameters as the agent gathers more information about the trainee. Based on the model, we formalize a method for adaptive goal setting for aerobic exercises which can be used by a coaching agent for personal as well as temporal adaptation. We show that our approach select goals that are rated highly by expert physical therapists as compared to other plausible goals.

Background

To improve overall cardiovascular health, the American Heart Association (AHA) recommends at least 150 minutes of moderate intensity aerobic exercise for adults. Aerobic exercise increases the overall energy expenditure through contraction of large muscle groups. In practice, the FITT-VP principle is used to prescribe exercises. The principle suggest that **F**requency (*f* days per week), **I**ntensity (*i* e.g., light, moderate, vigorous), **T**ime (*t* minutes per session), and **T**ype of exercise (in our case walking) can be used to modulate the **V**olume ($v = i \times t \times f$) of exercise for an individual (American College of Sports Medicine 2013). The volume per week correlates with energy expenditure in that time period. The intensity of an exercise can be represented in MET (metabolic equivalent) which is the ratio of the energy cost of performing it and the resting metabolic rate. Table 1 shows various walking exercises and valid values for duration and frequency that our coach uses for training.

The social coach is designed in a fashion similar to a symbolic, relational cognitive system (Langley, Laird, and Rogers 2009). Figure 1 shows the architecture employed to design the social coach. The coach has reasoning modules for two primary behaviors - interaction and goal setting that rely on information from each other for reasoning. Each of these modules has a short-term working memory that is encoded as a graph. It contains the coach's current

state which includes its hypotheses about and experiences with the trainee as well as its processing state. A procedural memory of rules drives the coach's reasoning. Whenever a rule's left-hand side condition matches the working memory, it fires and its actions change the working memory. If several rules match at the same time, numeric preferences are applied to force an order.

The coach has a specialized long-term memory that stores information from past interactions with the trainee as well as their performance on exercises recommended by the coach previously. This information is made available in the working memory to drive behavior. The coach also contains a reasoning controller. Its primary job is to determine which reasoning module the coach should employ - if it should run the interaction controller or the goal scheduler. It uses the current state of the interface as well as some pre-encoded heuristics to determine which reasoning module to run.

Interactive behavior is employed to measure a trainee's state, to gather more information about their performance, or to achieve joint agreement about goals. Interactions between the coach and the trainee are event-driven and may be initiated either by the trainee or the coach. Goal setting behavior looks at past interaction and performance data in the long-term memory and creates a schedule of walking goals.

Computational Formulation of Goal Setting

The goal setting theory (Locke and Latham 2002) provides evidence that to be maximally effective, the goals should be difficult yet attainable. For aerobic exercises, difficulty can be expressed as the volume prescribed given the trainee's physical state. The attainability can be considered as how safe the prescribed goal is for a given physical state and how likely is it to be successfully completed. While our representation does not explicitly describe these quantities, the proposed method and heuristics implicitly captures them. The theory also suggests that proximal goals are more ef-

fective than distal ones. This implies that the coach must generate a schedule that can be acted on by the trainee in the near future. Further, it recommends setting specific goals (e.g., *moderate intensity walk for 20 minutes*) as they are more effective than general ones (e.g., *exercise more*). Following these suggestions, the coach is designed to maintain a fully specified schedule of walking exercises for seven contiguous days starting on the current day. A trainee can look at this schedule in the smartphone application. On the days that the trainee is expected to exercise, the application shows the scheduled exercise (e.g., *brisk walk*) and its duration (e.g., *20 minutes*). Given this desiderata, the goal setting problem can be considered to be composed of two sub-problems: *weekly scheduling* - determining the exercise volume to be pursued in a week given a long-term goal and *daily scheduling* - i.e. distributing the weekly goal to specific days.

Weekly scheduling: Consider a trainee who has an aerobic capacity of c_0 at the beginning of the intervention and is advised to achieve the goal g_n for a healthy lifestyle. The weekly scheduling problem is to generate a schedule of exercise goals $G_{[1,n]} = \{g_1, \dots, g_n\}$ for weeks $[1, n]$. Following the FITT-VP principle, the weekly goal is represented as a tuple $g_w = (n, i, d, f)$ where n is the exercise name, i its intensity, d the duration of a session, and f sessions in a week. As the trainee achieves the exercise goal each week, their capability grows as a function of their prior capability and the exercise schedule $c_w = m(c_0, G_{[1,w-1]})$. A week's goal is g_w selected such that it requires the capability c_w for successful completion. This relationship is captured in a mapping function $g_w = r(c_w)$. Consequently, the goals get incrementally harder as the trainee's capability grows. At week n the trainee can achieve the long-term goal g_n which requires capability $c_n > c_0$.

Daily scheduling: Given goals for two consecutive weeks w and $w + 1$ and number of sessions completed in w , the daily scheduling problem at day d_w in week w is to select days in the interval $[d_w, d_w + 7]$ on which sessions will be scheduled where some days in $[d_w, d_w + 7]$ may be in week $w + 1$. This should be done so that the opportunity for the trainee to achieve their weekly goals is maximized and adequate rest days are scheduled between sessions.

Trainee Aerobic Capability Model

Given the concepts introduced in the previous section, an optimal weekly schedule can be computed using a standard forward search algorithm. However, few crucial challenges must be addressed. First, measuring a trainee's capacity c_0 through a mobile platform is neither straightforward nor precise. Experts rely on questions about a trainee's lifestyle as well as physical tests to estimate a trainee's capacity. To automate this process, several problems related to computer vision must be solved. Our coach uses a questionnaire about how active the trainee is in a typical week to generate an initial hypothesis about the trainee's aerobic capability. However, this assessment is error prone.

Second, as every trainee is different, the model m required for scheduling cannot be fully specified during design time. It has to be fit to every individual who the coach trains. Additionally, the coach's goals should be reasonable even at the

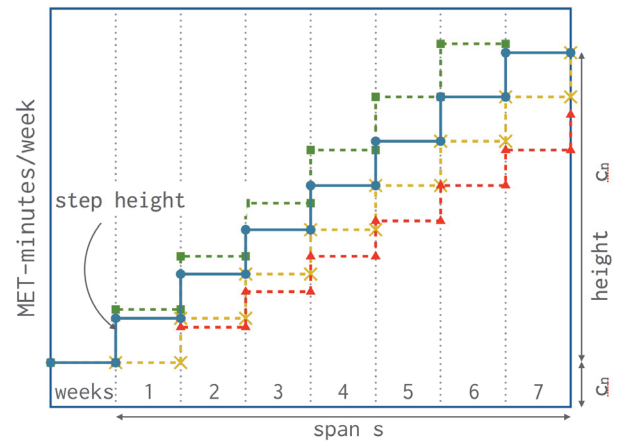


Figure 2: An adaptable model for aerobic capability. Blue line (●) represents the staircase model for growth in capability and green (■), yellow (×), and red (▲) its revisions.

beginning of the program. This is a requirement because if the goals are too hard a trainee might injure themselves or if they are too easy then the trainee might not be motivated to continue the coaching. Finally, an ideal solution for daily scheduling requires knowing how much time the person has available on each day and distributing the sessions accordingly. However, without knowledge of a trainee's schedule, it is hard to determine which days are ideal.

These challenges motivate an adaptive, knowledge-rich approach to designing the model as well as scheduling goals. The adaptivity of our method alleviates the errors in assessing a trainee as well as the lack of information about a trainee's schedule. While the initial goals may not be ideal due to assessment errors, the model can be refined based on observations made during training. To develop the parameterized model described below, we analyzed how experts prescribe exercises. The model encapsulates the structure experts rely on to prescribe exercises. The parameters can be revised online to fit different trainees.

We assume that people differ in two important ways: in their aerobic capability at the beginning of the program and in how quickly their physical capability can grow. The coach represents a trainee's aerobic capability c_w as the exercise volume they can achieve in a week w . Given the intensity i METs, the duration of a session d , and the number of sessions in a week f , their weekly aerobic capability is computed as $c_w = i \times d \times f$. Not only is this measure of physical capability standardized across various aerobic activities, it provides a direct mapping ($g_w = r(c_w)$) between the activity goals and the capability required to achieve those goals.

To capture weekly growth in aerobic capability, we employ a staircase function. The model assumes that a trainee's capability grows as a staircase function of equally spaced (1 week) steps of uniform height. Figure 2 shows some examples of such a model. The height of the function at week w captures the capability c_w in that week. The step height captures the coach's hypothesis about how quickly a trainee's capability can grow. The model can be abstracted as the tu-

ple (c_0, c_n, s, o) where c_0 is the height of the floor of the staircase, c_n the height of the highest step, s the span, and o is an offset. The model has two parameters that can be adapted by revising this tuple as follows:

1. *change-step*: The step height can be revised by increasing (or decreasing) the staircase's span (on the x-axis) by δ weeks making it easier (or more difficult). For example, increasing the step size of the blue staircase function (●) in Figure 2 by decreasing the span s by 1 week results in the green staircase function (■). This corresponds to a revision in the hypothesis about how quickly a trainee's capability can grow.
2. *shift*: The staircase function can be shifted forward on the x-axis by δ weeks. For example, shifting the blue staircase function (●) in Figure 2 by 1 week ($o = o + 1$) results in the yellow staircase function (×). This corresponds to a revision in the next week's capability without revising the hypothesis about capability growth.

The revisions can also be applied together. For example, the red staircase (▲) is achieved by shifting the blue one (●) by a week and increasing its span by a week.

Adaptive Goal Setting

Here we describe how the coach employs interactions and goal setting to coach a trainee toward the AHA goal.

Assessment: Before planning a schedule of exercise goals, the coach must assess a person's aerobic capability at the beginning of the program (c_0). This is critical for initializing the trainee model described earlier and consequently to schedule appropriate goals. Through a series of assessment questions, a person reports the duration and frequency of activity categories (with varying intensity) they undertake in a typical week. The coach aggregates these volumes into an assessment capability c_0 (shown in Figure 2). c_0 demarcates the floor (the interval before week 1) of the staircase function. Volume corresponding to the AHA goal (c_n) demarcates the height.

Then, the coach calculates a set of choices for the first week's goal by using the values specified in Figure 1. The coach selects the lowest intensity exercise (i) that can achieve the same volume as what is assessed. For this exercise, the coach determines the least duration (d) and corresponding frequency (f) will achieve the same volume as the assessed capability. These values i, d, f map a trainee's capability to the exercises in the program. The goal choices for the first week (g_1) are computed by incrementally adding 5 minutes to the duration until the maximum duration is reached. After compiling this set, the coach asks the trainee to pick a goal that they are most comfortable in attempting in the first week. The difference between c_0 and volume of the chosen goal determines the height of the first step, h (step height in Figure 2). Assuming uniform height steps, the span of the staircase is computed by $s = (c_n - c_0)/h$. This span is the projected time the trainee will take to reach their goal. People have a reliable estimate of how successful they can be at a task (Bandura 1994). Incorporating their choice in the model ensures that coach starts with a reasonable hypothesis that can be further refined online.

Daily Scheduling: At the beginning of the week w , the coach distributes the $g_w.f$ sessions between week days such that the rest days (days with no scheduled walking) are uniformly spaced. The trainee is expected to report their performance every day. They can report that they successfully performed the activity (*done*), that they tried the exercise but could not complete it (*almost*), and that they did not do the exercise (*nope*). For scheduling, *almost* and *nope* are equivalent. If successful, the original schedule is maintained. Otherwise, the coach redistributes $g_w.f$ in the remainder of week w such that rest days are uniformly spaced. As the trainee moves through week w , the coaching starts scheduling week $w + 1$ similarly to maintain a schedule for seven contiguous days. If the trainee is unable to achieve the daily goal, this rescheduling ensures that they get another chance to achieve it as long as there are enough days in the week.

When the person reports on an activity, the coach engages them in further interactions to gather more information about their performance. The trainee can report in three ways. On reporting *done*, the coach asks the person report how tired they were during the exercise on a 5-point Likert scale. Rating of 1 in our scale corresponds to no exertion ("Not tired at all"). 2 reflected very light exertion ("Little tired: breathing felt easy"), 3 moderate "Tired, but can still talk", 4 challenging ("Really tired: felt"), and 5 impossible ("So tired: had to stop"). The ideal exercise for the trainee should make them tired but they should still be able to talk (3 on the Likert scale). On reporting *nope* or *almost*, the coach asks the person to pick the reason why they did not do the activity. The reasons include if they didn't have time, don't enjoy it, don't find it useful, and found the activity too hard.

These interactions elicit information that is useful in revising the hypothesis if needed. The daily scheduling that occurs during the week can be considered as an evidence-collection phase in which the coach is observing the trainee's performance on a prescribed goal. This evidence is incorporated in the coach's reasoning as below.

Weekly Scheduling: Based on the observations and interactions in week w , the coach may update the staircase model for weeks $> w$. This update is triggered by observations that deviate from expected performance of the person on the selected goal g_w . If in week w , the determined activity goal g_w is appropriate for a person's capability c_w , it is expected that the person can achieve it with an average exertion of 3. Any diversion from this should trigger a revision. The coach first diagnoses if the current model under- or over-estimates the person's capability and then adapts it as follows.

- *regress* revision: The coach makes a regress revision if,
 - the person is unable to complete at least 50% of g_w or,
 - the person completes $> 50\%$ of g_w but reports an average exertion ≥ 4 or there exist at least one report in which the reason for not completing an exercise is that it is *too hard*

If these criteria is met, the coach assumes that the staircase model is not only overestimating the trainee's capability this week but it also overestimates how quickly the capability can grow. The coach revises the model by increasing the span of the staircase (a *change-step* manipulation with $\delta = 1$) and thereby making the staircase less steep. This

represents a revision in the coach’s hypothesis about how quickly the person’s capability can grow. The coach also shifts the staircase function by $\delta = 1$. As the person failed g_w , this revision ensures that the next week’s goal will be easier than this week.

- *progress* revision: The progress revision occurs when the coach observes that the person completed at least 75% of g_w and reported an average exertion of ≤ 2 . This signals that the goal scheduled given the model’s estimate of the person’s capability and growth is too easy. Therefore, the model should be revised to reflect a faster growth in capability. The coach revises the model by decreasing the span of the staircase (a *change-step* manipulation with $\delta = -1$) and thereby making the staircase harder.
- *shift* revision: The shift revision occurs when the person’s non-performance is caused by something other than their aerobic capability. The coach makes this determination if the person completes 50% – 75% of g_w , the criterion for a regress revision was not met, and there exist at least one report in which the reason for not completing the daily goal is that the person was *too busy*. This suggests that the g_w may be appropriate and if given an opportunity again, the person may be able to achieve it. The coach shifts the staircase function by $\delta = 1$ week to give the person another opportunity at completing the goal without revising the hypothesis about growth in their capability. There may be several other reasons for why a person may not achieve a goal lack of motivation to invest resources or no expectation of internal/external reward. Future variations of the coach will address diagnoses of these reasons and strategies to overcome them.

Weekly goal setting: Given a person’s capability c_w from the model, the coach computes the activity goal g_w as follows. First, the coach generates possible combinations of intensity, duration, and frequency. For every activity of intensity i under consideration and for every frequency value f , the coach computes the relevant duration value $d = c_w / (i \times f)$ approximated to the closest multiple of five. Any combination in which the duration is higher than the maximum or lower than the minimum is rejected. If $w = 1$, the combination that matches the person’s choice (in assessment) is set as the goal. For $w \neq 1$ the following filters are applied incrementally:

- if $c_w = c_{w-1}$, the combination matching the previous week’s goal is selected as this week’s goal.
- if $c_w < c_{w-1}$ it implies that the previous week’s goal was harder than what the person could achieve. Therefore, the combinations that are harder than the previous week’s goal $((g_w.i > g_{w-1}.i) \vee (g_w.i = g_{w-1}.i \wedge g_w.d > g_{w-1}.d) \vee (g_w.i = g_{w-1}.i \wedge g_w.d = g_{w-1}.d \wedge g_w.f > g_{w-1}.f))$ are rejected. From the remaining combinations, the easiest combination is selected to be this week’s goal. Selecting the easiest goal ensures that a relatively safer goal is scheduled given the constraints derived from the model.
- if $c_w > c_{w-1}$, it implies that the person can attempt a goal that is harder or at least equal to the previous week’s goal. The combinations that easier than previous week’s goal $((g_w.i < g_{w-1}.i) \vee (g_w.i = g_{w-1}.i \wedge g_w.d < g_{w-1}.d) \vee$

$(g_w.i = g_{w-1}.i \wedge g_w.d = g_{w-1}.d \wedge g_w.f < g_{w-1}.f))$ are rejected. As earlier, the easiest combination is selected from the remaining combinations.

Evaluation

We are interested in the following desiderata of a coaching agent: *adaptivity* - the coach is able to adapt the sequence of goals to different trainees as well as to how a trainee’s capability changes when they begin to exercise and *effectiveness* - the coach can prescribe goals that are difficult yet attainable, useful, proximal, and specific. Below, we present our analysis of the proposed approach on these desiderata. The experiment was conducted by simulating various trainee profiles and observing the behavior of the coach. Given our domain analysis that trainees differ along two dimensions, we simulated 6 profiles with 3 levels of initial capability — no, low, and moderate and 2 levels of capability growth profiles — low and moderate. We did not simulate high growth profiles as this is not characteristic of our target population that comprises of sedentary, overweight trainees.

Adaptivity: The coach’s adaptivity is demonstrated in Figure 3 which shows the goals scheduled by the coach for three different trainee profiles for 8 weeks. The blue bars represent the goals for a trainee profile who did not exercise before the training and whose aerobic capability grew moderately. We see that the goals became incrementally harder (increase in volume) as time progresses. This trajectory can be compared to the trainee (green bars) who did not exercise before the coaching and whose aerobic capability grows slowly. By week 4 the goal increases to 25 minutes of moderate walk 5 times a week. However, this trainee is unable to achieve this goal successfully. We can see that until week 4, both trainees are given similar goals. However, week 5 onward the second trainee is given easier goals as the coach revises its hypotheses about the trainee at week 5 given observed performance until week 4. The third trainee (yellow bars) began at harder goals than the previous two as they assess higher. But, their capability grows similarly to the second trainee. Therefore, the coach makes the goals easier at week 5. These results show that the coach can adapt the goals to changes in a trainee’s capability (temporal adaptation) as well as to different trainees (personal adaptation).

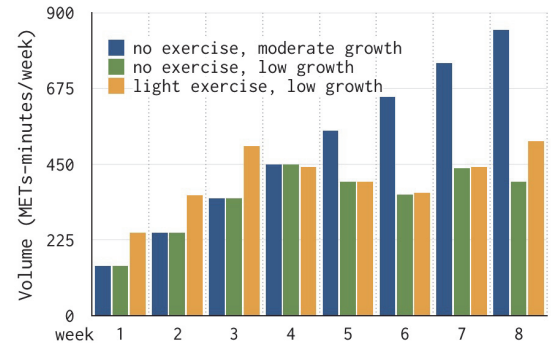


Figure 3: Prescribed exercise volume for 8 weeks for 3 profiles.

	Safe		Useful	
	$\chi^2(1) = 76.56, p < 2.2e^{-16}$		$\chi^2(1) = 40.21, p = 2.267e^{-10}$	
	Agree	Disagree	Agree	Disagree
Coach	266	19	266	22
Controls	379	197	425	151
	Likely		Difficult	
	$\chi^2(1) = 79.78, p < 2.2e^{-16}$		$\chi^2(1) = 23.32, p = 1.37e^{-06}$	
	Agree	Disagree	Difficult	Easy
Coach	275	13	156	132
Controls	393	183	409	167

Table 1: Contingency tables for binary expert ratings of coach-selected v/s control goals.

Effectiveness: The coach should set goals that are difficult yet attainable, useful, proximal, and specific. The latter two desiderata are achieved by design as described in the previous sections. To evaluate the earlier two, we asked 6 physical therapists to judge the coach’s performance. Each trainee profile was described to experts in terms of initial assessment, desired long-term AHA goal, history of weekly adherence, and average exertion scores. Experts were asked to judge coach-selected goals relative to 2 control goals picked by a different expert such that they are reasonable and plausible given past history of trainee performance. Experts rated safety, usefulness, likelihood of completion and difficulty of all 3 goals on a 6-point Likert scale (see Figure 4) each week for the 8-week program for all 6 profiles. The neutral option on the Likert scale was deliberately omitted to force judgment in a direction. Importantly, the experts participating in the study were blinded to the fact that these goal recommendations originated from an algorithm to reduce/avoid any unintentional biases. They rated these goals under the knowledge that they were prescribed by other experts such as themselves.

As seen in Figure 4, experts rated the coach-selected goal higher on safety, attainability, and usefulness compared to control goals for all 6 profiles across 8 weeks. These goals were also rated easier than controls. To test if these differences were statistically significant, we compared expert ratings converted to a contingency table (see table in Figure 4). For this, the 6-point Likert scale levels were collapsed to obtain a binary classification (e.g., agree versus disagree for safety, attainability, and usefulness and goal difficulty rated as difficult versus easy). Further, the 2 control goals were merged and compared to the coach-selected goals. As shown, χ^2 -squared tests revealed that these differences were statistically significant ($p < 0.001$).

Our results suggest that the coach selects goals that have a higher chance of being safer, are more likely to be completed, and useful to achieve a defined long-term goal as compared to other plausible goals. Further, our method selects goals of appropriate difficulty given history of performance (see box plot *d* in Figure 4). It shows that the goals preferred by the coach are more likely to fall between *somewhat easy* and *somewhat difficult*. This demonstrates that the goals selected by the coach are reasonable for various

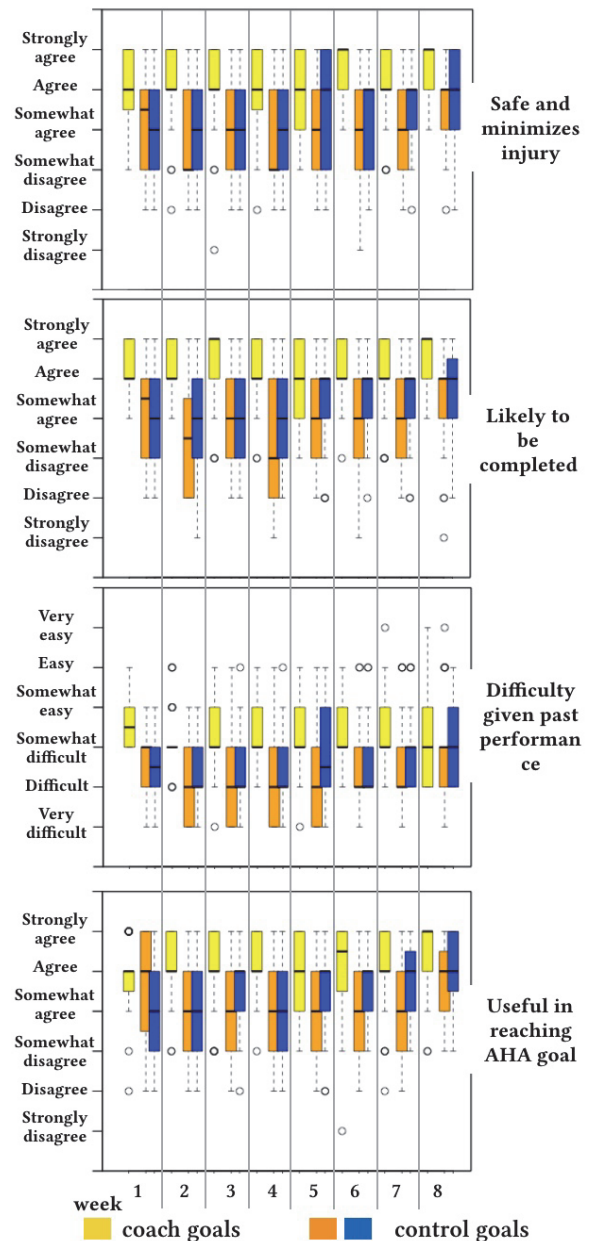


Figure 4: Box plots showing expert ratings of weekly goals for an 8 week program. Boxes depict inter-quartile range, horizontal bar medians, dashed lines range, and circles outliers.

trainee profiles and have a higher chance of being successfully achieved than other comparable goals.

Related Work

Past work in AI on designing social agents has largely looked at designing personal assistants for knowledge tasks (Myers et al. 2007). In the past decade, technology for health has gained some attention in AI research, however, use of intelligent reasoning and decision is in its infancy. Re-

search has focused on mechanisms for intelligent reminding in which for an expert-picked goal (e.g., *make a diary entry*), an AI system generates personalized reminders (Buttussi and Chittaro 2008). Others (Lisetti and Wagner 2008) address the design of conversational agents or dialog systems that conduct motivational interviewing to influence adherence to a behavior change intervention. These approaches rely on expert coaches to design relevant goals while the role of the agent is limited to delivering these goals to the trainee. We take a step beyond these and describe a method through which an intelligent agent can generate personalized exercise programs for trainees.

In order to develop our approach, we proposed a trainee model that drives the coach's expectations about the trainee and is useful in picking ideal goals. Previously, the trainee (or learner) models have been studied by the intelligent tutoring systems community (Desmarais and d Baker 2012). These models assume that what drives a learner's performance is a set of discrete skills that they possess. This discrete representation is not sufficient for representing factors that influence a trainee's performance on exercises such as walking. These models usually represent beliefs about if the learner knows cognitive skills such as addition or multiplication. This is not sufficient for coaching exercises. Even if a trainee knows how to walk, they can have substantially different performance on walking for 15 minutes versus 30 minutes. Moreover, the models investigated are static and are learned a priori. Our work develops a new kind of a predictive model that is targeted toward representing physical skills and capability required for walking. The model can be revised online and gradually adapts to each specific trainee.

Conclusions and Future Work

In this paper, we proposed a design of an interactive agent that has several properties desirable in a coach including being able to adapt the training to different individuals and being able to select goals that are appropriate for a trainee given past history of performance. However, there are a few shortcomings to our approach. A primary limitation is that the model cannot be adapted such that the predicted capability in a week $w \geq 1$ is lower than the initial capability measurement c_0 . Therefore, if initial capability is overestimated to a high degree, the coach may never recover from the error. Similarly, the staircase model is constrained to be adapted by at most a week, ($\delta = 1$) despite the method allowing for any arbitrary number, precluding quicker adaptations. These are empirical questions that can be answered by deploying the coach in the target population in the future. A limitation of the study is that experts were not probed about the rationale underlying their judgments. Insights so derived will inform formulating a more expressive model of growth in aerobic capability as well as elicit heuristics that experts employ to adapt exercises. The results presented here provide evidence that our methods can be effective. To further strengthen the claims about the efficacy of our approach, we want to deploy and study the coach working with real people in future. Finally, we need more work to extend the coach design to incorporate reasoning about the influence of other cognitive and affective processes on health behaviors.

Acknowledgment

This research was supported in part by funding from Xerox corporation and by NSF-SCH grant 1346066. The authors thank the anonymous reviewers for their insightful comments and suggestions.

References

- American College of Sports Medicine. 2013. *ACSM's guidelines for exercise testing and prescription*. Lippincott Williams & Wilkins.
- Bandura, A. 1994. *Self-efficacy*. Wiley Online Library.
- Buttussi, F., and Chittaro, L. 2008. Mopet: A context-aware and user-adaptive wearable system for fitness training. *Artificial Intelligence in Medicine*.
- Desmarais, M. C., and d Baker, R. S. 2012. A review of recent advances in learner and skill modeling in intelligent learning environments. *User Modeling and User-Adapted Interaction*.
- Garber, C. E.; Blissmer, B.; Deschenes, M. R.; Franklin, B. A.; Lamonte, M. J.; Lee, I.-M.; Nieman, D. C.; and Swain, D. P. 2011. Quantity and quality of exercise for developing and maintaining cardiorespiratory, musculoskeletal, and neuromotor fitness in apparently healthy adults: guidance for prescribing exercise. *Medicine & Science in Sports & Exercise*.
- Langley, P.; Laird, J. E.; and Rogers, S. 2009. Cognitive architectures: Research issues and challenges. *Cognitive Systems Research*.
- Lisetti, C. L., and Wagner, E. 2008. Mental health promotion with animated characters: exploring issues and potential. In *AAAI Spring Symposium: Emotion, Personality, and Social Behavior*.
- Locke, E. A., and Latham, G. P. 2002. Building a practically useful theory of goal setting and task motivation: A 35-year odyssey. *American Psychologist*.
- Myers, K.; Berry, P.; Blythe, J.; Conley, K.; Gervasio, M.; McGuinness, D. L.; Morley, D.; Pfeffer, A.; Pollack, M.; and Tambe, M. 2007. An intelligent personal assistant for task and time management. *AI Magazine*.
- Shilts, M. K.; Horowitz, M.; and Townsend, M. S. 2004. Goal setting as a strategy for dietary and physical activity behavior change: a review of the literature. *American Journal of Health Promotion*.