Smart Health 28 (2023) 100389

Contents lists available at ScienceDirect

Smart Health

journal homepage: www.elsevier.com/locate/smhl

Towards Dynamic Action Planning with user preferences in Automated Health Coaching

Ajith Vemuri [a,*], Megan Heintzelman [b], Alex Waad [b], Matthew Louis Mauriello [a], Keith Decker [a], Gregory Dominick [b]

[a] *Computer and Information Sciences, University of Delaware, Newark, DE, USA*
[b] *Behavioral Health and Nutrition, University of Delaware, Newark, DE, USA*

ARTICLE INFO

Keywords:
Virtual coach
Wearable Devices
Interventions
Physical Activity
Individualization
Planning/Scheduling
Preferences
Self-Determination Theory

ABSTRACT

Health coaching is an evidence-based approach to help individuals adopt health behaviors such as physical activity (PA). However, human health coaching is limited by a lack of real-time data and scalability. In this paper, we present the architecture and functionality of a novel semi-automated health coach that we call BeSMART that automates the goal modification and dialogue functionalities of human health coaching using a commercial health tracker and mobile phone. Results from an acceptability study (N = 10) indicate that users have difficulty fully completing all actions as initially planned. To build planning systems that can dynamically adapt action plans to such changes, we show that user preferences are vital, and that users find re-planning acceptable. Towards solving this problem, we extracted and categorized user preferences from health coaching sessions and propose extensions to Hierarchical Task Network (HTN) representations. When taking users personally re-planned activities into account, post-hoc analysis shows that users are much more likely to complete their overall PA goals.

1. Introduction

Health coaching (Oddone et al., 2018; Wolever et al., 2013), a personalized guidance technique, has been shown to be an effective method to change physical activity behavior. What sets health coaching apart from other strategies is that it is client-centered. Coaches and clients work together to identify personalized health goals that the client wants to accomplish. Through conversations, health coaches help clients create goals that are Specific, Measurable, Attainable, Realistic, and Time-bound (SMART). For example, "Over the next 7 days, I will walk at a moderate intensity for 30 min, 3 times a week". In addition to SMART goals, the health coach and client will negotiate plans that detail the strategies to actualize the goal (i.e. *action plan*) as well as for overcoming barriers to their goal (i.e., *coping plan*). Within the context of the above SMART goal, an action plan would outline the specific days, times, and locations where the client intends to walk. A coping plan would then outline the specific barriers that the client identifies as most pertinent to them and realistic strategies for overcoming them. Although effective, health coaching has limitations including the difficulty to scale given its highly individualistic approach, and the lack of integrated objective data collected in real-time.

One potential solution to the identified limitations of human health coaching relies on the ubiquity of smart phones and wearable devices. These technologies provide a rich sensor platform and medium for both collecting data and delivering scalable behavioral interventions to influence health outcomes. However, leveraging such platforms and automating the coaching process has mostly been limited to theoretical studies (Márquez-Hernández et al., 2022) or studies of systems with limited interactivity (Althoff, Clark, & Leskovec, 2016; Gupta et al., 2020; Pérez-Rosas et al., 2018).

* Corresponding author.
E-mail address: kumar@udel.edu (A. Vemuri).

https://doi.org/10.1016/j.smhl.2023.100389
Received 14 March 2023; Accepted 15 March 2023
Available online 23 March 2023

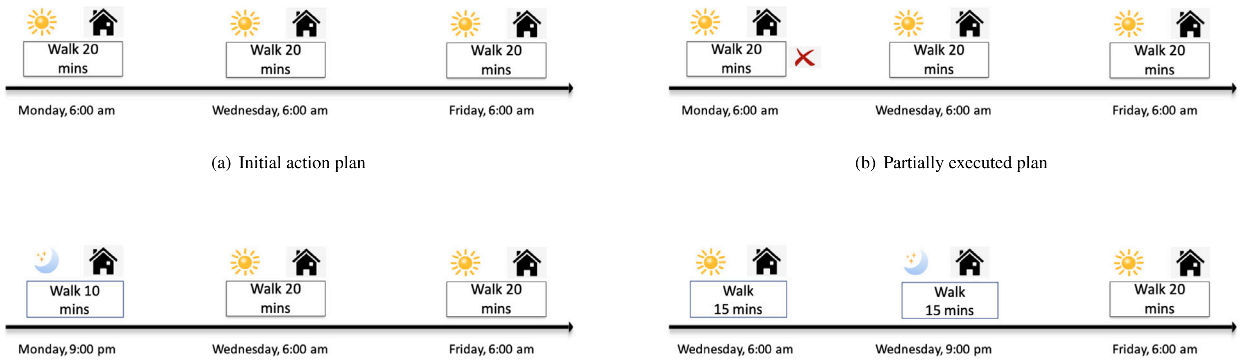

(a) Initial action plan

(b) Partially executed plan

(c) Recommendation-1 for new action plan

(d) Recommendation-2 for new action plan

Fig. 1. Dynamic action planning scenario. The system will adapt to changing contexts of users (b) and generates recommendations (c) and (d) based on user preferences.

In this paper, we present the following contributions to automate the health coaching problem:

1. A novel semi-automated coaching system, BeSMART, that helps users to achieve and modify their PA goals via text message dialogues. The modular design of BeSMART enables additional enhancements making it an ideal platform for similar automated coaching systems.
2. From our pilot data, we show that goal achievement is an iterative process where users partially succeed and fail multiple times in the process of changing their PA behavior. Based on this observation, we introduce a dynamic action planning problem, that when solved, could enhance the effectiveness of health coaching.
3. From our analysis of transcript data of conversations between certified human health coaches and participants, we show that a comprehensive representation of preferences over planning goals, sub-goals, and primitive actions is needed to solve the dynamic action planning problem that the current literature lacks.
4. We propose theoretical extensions to existing planning and scheduling frameworks that will enable development of new dynamic planning and scheduling algorithms that can adapt to changing user contexts based on their individual preferences.

The BeSMART system is a personalized coach that automates most of the human health coaching aspects by leveraging FitBit devices and smart phones. The BeSMART system achieves individualization, privacy, and scalability by leveraging a multi-agent architecture and deploying a personal coach (agent) for each user. BeSMART holds varied dialogues with users throughout the week to remind, motivate, track, and modify a user's physical activity goals. By switching between multiple finite state machines, BeSMART holds semi-customized conversations with users based on their weekly PA goal progress. To test the acceptability of BeSMART (Heintzelman, Dominick, Vemuri, & Decker, 2022), we conducted a 6 week trial with 10 study-eligible participants. Participants first met with a human health coach to establish their initial PA goals as well as their action and coping plans. Baseline measurements of physical activity and blood pressure occur in the first week, followed by two weeks of the BeSMART system, a mid-point human coaching session (week 3), two more weeks of BeSMART, and a final measurement week and exit surveys (see Section 5).

The results of the acceptability trial show that participants who try to change their PA behavior achieve their PA goals over multiple attempts with some successes, partial executions, and failures. The data shows that the action plans executed by the participants are dynamic and diverge from the initial action plans set with the human health coaches. We believe that to further improve the effectiveness of automated health coaching, digital coaches should dynamically adapt to a user's actual current goal fulfillment, and produce both reminders and activity recommendations accordingly. For example, Fig. 1(a) shows a user's action plan to achieve their PA goal of 60 min, by walking 20 min three times a week on Mondays (6:00 am, location: home, weather: sunny), Wednesdays (6:00 am, location: home, weather: sunny) and Fridays (6:00 am, location: home, weather: sunny). Come Monday, as shown in Fig. 1(b), the user may only partially fulfill their goal (i.e., walking for 10 min instead of 20). In cases like these, automated coaches may adapt by either recommending an alternative time such as Monday at 9:00 pm as shown in Fig. 1(c), or increasing Wednesday's goal from 20 to 30 (shown as two 15 min bouts) as in Fig. 1(d).

These recommendations cannot be made arbitrarily, but should be based on user's preferences of when and how they intend to achieve their PA goals. From an analysis of transcripts from conversations with health coaches, we identified, extracted, and categorized types of preferences such as conditional preferences over activity types, preferences over how to break up goals, and preferences over doing activities with other people, all of which are needed to make such recommendations. We also propose extensions needed for plans generated by traditional Hierarchical Task Network (HTN) planners to develop new algorithms which can make such recommendations while maximizing user preferences.

2. Related work

Here we discuss background on Self-Determination Theory as it informs the design of our Automated Health Coach as well as highlight our focus on preferences in planning/scheduling.

2.1. Self-determination theory

Self-Determination Theory (SDT) provides a strong theoretical framework for understanding human motivation (Kinnafick, Thøgersen-Ntoumani, & Duda, 2014; Ryan, Patrick, Deci, & Williams, 2008; Teixeira, Carraça, Markland, Silva, & Ryan, 2012; Wilson, Mack, & Grattan, 2008), and has significant empirical support regarding its utility for developing PA interventions, and predicting PA adoption and maintenance. SDT posits that motivation occurs along a continuum from non-autonomous to completely autonomous. In a PA context, autonomous motivation is achieved when PA is done for personal enjoyment, to attain valued goals, or to improve current health. According to SDT, three basic psychological needs foster autonomous motivation; behavioral regulation (e.g. goal-setting, self-monitoring, and cues to action), perceived competence (i.e., self-efficacy), and relatedness (e.g. emotional connection with others). Moreover, the Transtheoretical Model (i.e. stages of change) and Change Processes are proven effective for impacting behavior change (Fahrenwald & Walker, 2003; Marcus et al., 1998; Marcus, Rakowski, & Rossi, 1992), and are recommended to inform the development of tailored behavior change messages. The BeSMART coach applies Self-Determination Theory and the Transtheoretical Model to address personal motivation and relevance to the context of the unique participant.

2.2. Automated health coaching

Work exists that attempts to apply behavioral theory to engage in a personalized and goal-oriented interactive dialogue to motivate clients towards behavior change (Dennis et al., 2013; Eakin, Lawler, Vandelanotte, & Owen, 2007; Kivelä, Elo, Kyngäs, & Kääriäinen, 2014; Mahon et al., 2018; Oddone et al., 2018). Many dialogue-based health systems are limited to classification or summarization rather than dialogue interaction (Althoff et al., 2016; Gupta et al., 2020, 2021; Pérez-Rosas et al., 2018). Corpus-based chatbots exist, but they lack the structure and controllability necessary for a successful health intervention that avoids harm (Bickmore et al., 2018; Reiter, 2020). Additionally, to our knowledge, most existing non-theoretical dialogue system implementations are limited on tailoring (Bickmore, Schulman, & Sidner, 2013; Düking et al., 2020; Muntaner, Vidal-Conti, & Palou, 2016; op den Akker, Cabrita, op den Akker, Jones, & Hermens, 2015; op den Akker, Jones, & Hermens, 2014; Rupp, Michaelis, McConnell, & Smither, 2018; Svetkey et al., 2015), which is shown to increase user engagement and motivation (Mitchell et al., 2021; op den Akker et al., 2014). Moreover, few coaching platforms help users establish SMART health behavior goals nor do they include detailed action and coping plans.

2.3. Preferences in planning/scheduling

Prior work exists that accommodates preferences into planning algorithms (Brafman & Chernyavsky, 2005; Jorge, McIlraith, et al., 2008; Shaparau, Pistore, & Traverso, 2006; Sohrabi, Baier, & McIlraith, 2009). These approaches generate plans by choosing the most preferred decomposition, however they lack a comprehensive preference representation like conditional preferences over actions (Boutilier, Bacchus, & Brafman, 2013; Boutilier, Brafman, Domshlak, Hoos, & Poole, 2004; Brafman & Engel, 2009; Gonzales & Perny, 2004), or preferences over coordination relationships between goals or actions, which enable dynamic planning over preferences. For dynamic environments such as physical activity domains, planning, scheduling and execution are tightly intertwined (Decker, 1996; Lesser, 2002; Smith, Gallagher, & Zimmerman, 2007; Wagner & Lesser, 2000). Automated coaches should be able to switch between alternate plan decompositions seamlessly, requiring scheduling systems to reschedule to adapt to changing contexts of users based on their preferences. Current work represents preferences in scheduling as temporal constraint problems (Khatib, Morris, Morris, & Rossi, 2001; Peintner & Pollack, 2004; Rossi, Van Beek, & Walsh, 2006). These scheduling algorithms can only generate schedules that maximize temporal preferences, but lack contextual preferences.

3. The BeSMART coach

The primary objective of BeSMART (bio-Behavioral System to Motivate And Reinforce hearT health) is to increase weekly minutes in moderate intensity PA among mid-life adults with hypertension via an automated coaching system that incorporates data collected from a Fitbit Charge 3 and a Withings blood pressure monitor. As a secondary objective, this study also explores how changes in PA impact blood pressure in this population. To achieve the above objectives, BeSMART uses a FitBit to track PA, a Withings Blood Pressure device to collect BP data, and Twillio to hold dialogues with users via SMS (text messaging). Each user's weekly PA goals are customized to the user's heart rate. The BeSMART system also allows users to modify their weekly goals, which is one of the differentiating factors from other similar studies that try to increase PA. In the sections that follow, we describe how these objectives are achieved.

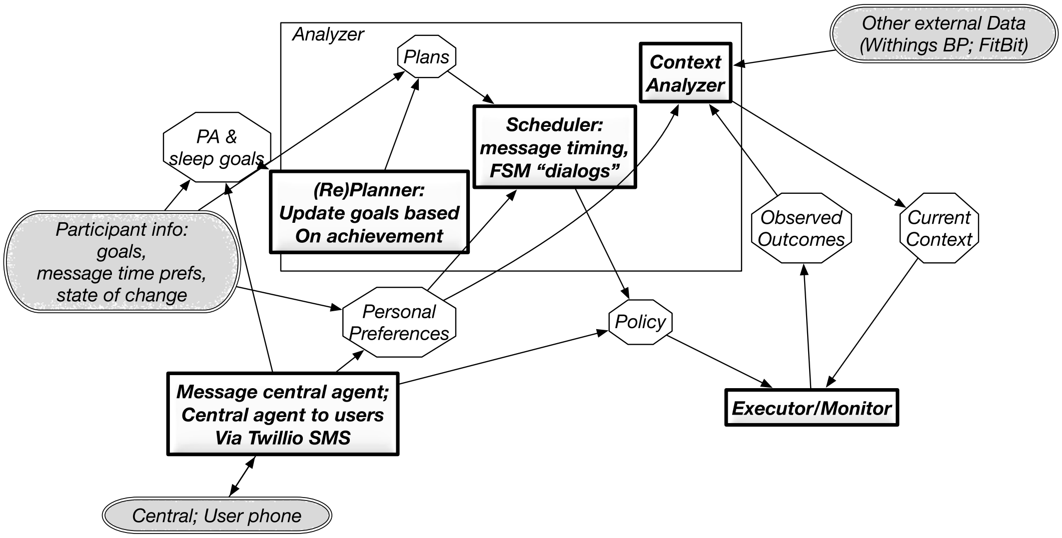

Fig. 2. BeSMART architecture.

3.1. Human coaching

Every user has a coaching session with a human health coach at the start of the study, as well as a midpoint conversation after study week 3. Human coaches partner with users to setup their initial PA goals as well as layout action and coping plans to achieve them. The PA goals are established using the Frequency, Intensity, Time and Type (FITT) principle. *Frequency* translates to the number of times the user plans to do PA in a week (the number of bouts). *Intensity* is measured using the Target Heart Rates (THRs) based on the fitness levels of individuals. Exercise physiologists calculate THRs for each individual and pass it on to the health coaches. *Time* represents the duration of PA the user plans to perform in a given week. *Type* represents the activity type the user intends to perform. A user can achieve a PA goal with more than one activity type.

Using the FITT principle, users choose a PA goal that they would like to start with. An example of a weekly PA goal in this study is as follows: Complete 1 bout of 20 continuous minutes of activity keeping your heart rate between 100–115 bpm on at least 3 days this week. The initial set of PA goals from which users can choose are recommend by exercise physiologists. After an initial goal is chosen by user, the user along with the human coach comes up with action and coping plans to achieve these goals. Activity plans consist of what days and times of the week users intend to exercise, and with whom. Coping plans consist of a list of barriers users might face while executing the action plan and a corresponding list of strategies to overcome these barriers. We will elaborate further in Section 7, on how the human coaching translates to preference requirements in planning/scheduling domains and offer solutions to automate this process.

3.2. BeSMART architecture

After the initial human coaching session and FITT-based action planning, the automated coach takes over. BeSMART monitors user progress towards their goals, and interacts with the user at three set times during the week. At the start of the week, BeSMART reminds the user of their weekly goal and current action plan. At midweek, BeSMART presents a snapshot of the user’s progress so far. At the end of the week, BeSMART reviews that week’s achievements, and converses with the user over whether to increase or decrease the user’s weekly goals and corresponding action plans, based on their success in the previous week. We discuss these components in more detail below.

BeSMART collects and analyzes wearable data (FitBit and Withings) to track user progress and conduct dialogues to help the user achieve PA goals. The architectural design of BeSMART, a derivative of SMART-ALEC (Vemuri, Decker, Saponaro, & Dominick, 2021) shown in Fig. 2 was not just influenced by the functionalities required by the automated coach in this study. As health coaching requires data privacy, scalability, and individualization, we opted for a multi-agent system (one agent per user) as an underlying platform over which BeSMART is built. As each agent acts as a private coach that can be deployed in a distributed fashion, this architecture can be scaled to very large numbers while maintaining privacy. Having a multi-agent system also helps with distributed coordination which we later show is an important feature of automated coaches.

The other distinctive feature of this design is its modularity. Separating functionalities into different modules allows easier integration of new features or removal of existing ones. For instance, the context analyzer module is responsible for reading and analyzing external data from wearables. Changing the wearable from FitBit to Apple Watch should not affect the rest of the modules in the architecture. Alternatively, if the automated health coach wants to implement a different behavioral paradigm like Just In Time Adaptive Intervention (Lentferink et al., 2017; Liao et al., 2018; Saponaro, Vemuri, Dominick, & Decker, 2021; Spruijt-Metz, Wen, & O'Reilly, 2015), adding in a new learner module without changing the other modules should be enough. Fig. 2 shows the agent-level architecture for each BeSMART coach. This architecture comprises four modules: Context Analyzer, Planner, Scheduler, and Communicator. Each of these modules are tightly coupled, dynamically interacting with each other during execution.

3.2.1. Context analyzer

The context or state of a BeSMART user is defined by wearable adherence on a daily basis, PA goal achievement on weekly basis, and heart rate, blood pressure measures being confined within safety thresholds. Variations in these states trigger interactions with the scheduler module which in turn sends messages to users. Users are expected to wear their FitBit throughout the day and take blood pressure readings once a day. Failure to take these measurements will result in reminders being generated and sent through text messages. PA goal progress is updated at the end of each day at a preset time 11:59 PM (US Eastern timezone). This is done by processing minute-level heart rate data of users to accumulate the total number of minutes spent in THR. Processed PA goals are depicted graphically to better help communicate with users. Different PA goal achievements at different points of the week lead to different types of dialogues held by BeSMART. For instance, a mid-week affirmation message is sent to someone who is on track to fulfill his/her goal vs an encouraging message to someone who is not. As BeSMART PA goals are exercise prescriptions, the system keeps a close watch on whether heart rates and blood pressures are exceeding safety levels, generating alerts to research coordinators in case they exceed certain thresholds (heart rate goes above 60% of a person's heart rate reserve; blood pressure > 160/90 mmHg). This allows research coordinators to personally contact the users and inform them about these events.

3.2.2. Planner

Users are given the option to modify the difficulty of goals on a weekly basis; goal modification takes the form of either increasing or decreasing difficulty. The planner presents a set of goal options for the upcoming week conditioned on the current week's goal achievement. These options are generated from a fixed set of goal configurations. A full goal achievement will generate options of either keeping the same goal or making it more challenging. A more challenging goal is either an increase in duration or intensity. Alternatively, a goal achievement of 50%–99% generates options of making the goal more challenging, keeping it the same and making the goal easier. Whereas a goal achievement of anything less than 50% will generate options of keeping the same goal for the upcoming week or making it easier.

3.2.3. Scheduler

BeSMART holds both two and one-way dialogues with users. As users need to be receptive to these dialogues, user preferences of receptivity are taken beforehand. Three preferred times for each user are input into the scheduler: start of the week (reminder) messages which are sent on Monday mornings, mid week (goal tracking) messages sent on Thursday evenings, and end-of-week (goal modification) messages sent on Sunday evenings. Because the end of week (goal modification) messages are important, the scheduler reschedules the dialogue on the same day at a different time if it detects non-responsiveness from the user. The user is asked for their preferred later time to continue the dialog. *This paper lays out the foundation to a similar idea but with rescheduling of activities rather than conversations.* If the user misses a scheduled activity, can the scheduler recommend an alternative preferred (activity,time,place) to achieve the user's goal?

3.2.4. Communicator

Connectivity among agents in a multi-agent system enables communication and coordination. The communicator module establishes asynchronous TCP connections among all the agents including a central agent. As communication is through Twillio via text messages, each agent should ideally have its own Twillio account allowing the compiled policy (FSMs) from the scheduler module to be deployed on to the executor module (executor here is the user smartphone). However, for the acceptability study, only one Twillio account was created and associated to a central agent to minimize cost. As a result, every agent holds its dialogues through the central agent. The central agent acts as a gateway to all the messages received and sent to users by other agents in the system. The central agent also hosts a web server (Flask) acting as a dashboard to researchers interacting with the agents. To maintain data modularity and privacy, finer level data such as minute-level heart rate and sleep data is held only at the agent level, stored in a dedicated private database. The central agent has access to higher level data including processed PA and dialogue text messages. This communication setup among agents will help implement additional features like coordination and transfer learning among different automated health coaches in future studies.

3.3. BeSMART dialogues

A key aspect of human health coaching is its personalized conversations. Though the messages of the BeSMART dialogues are pre-scripted, it achieves a good degree of customization across users by taking into account differences in: past and present PA, receptiveness of threatening health information, personal details, preferences, and responsiveness to the coach. BeSMART categorizes users based on their PA stage of change (Marcus et al., 1992), each user is either in the "preparation" or "contemplation" or the "action" stage of change. The current goal achievement is combined with the stages of change to make a more appropriate suggestion to increase PA. To motivate users better, especially the ones who are not on track to achieve their weekly goals, BeSMART informs users about the health impacts of sedentary behavior. Users are categorized either as monitors or blunters (Miller, Brody, & Summerton, 1988). *Monitors* seek out and monitor threatening or stressful information, but *blunters* tend to distract themselves to blunt the information's psychological impact. BeSMART sends motivational messages differently to monitors/blunters. BeSMART

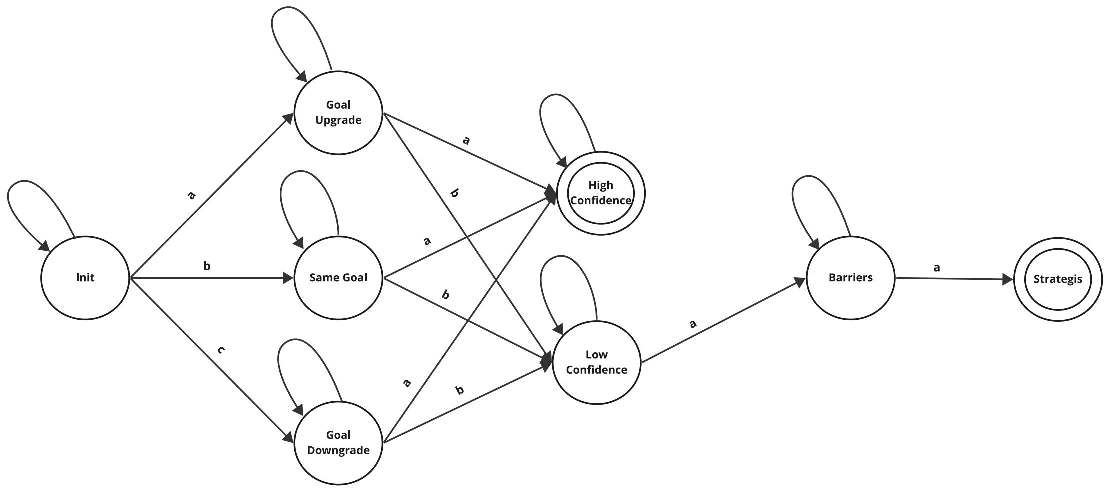

Fig. 3. BeSMART FSM for 50%–99% goal accomplishment.

also customizes messages based on the user's name, preferences over when to receive messages, adherence to wearables, and their responsiveness (or lack of) to messages.

BeSMART messages can broadly be classified into wearable reminders, alerts, goal reminder messages, goal assessment messages, and goal modification messages. Wearable reminders are messages to remind users to put on their wearables. These are sent early in the morning if the system detects lack of data from the previous day. Alerts generated for BeSMART administrators go out when the users do not respond or when user's heart rate or blood pressure measures are above recommended thresholds set by the American College of Sports Medicine. Goal reminder messages are always one-way dialogues and end-of-week messages are always two-way dialogues. Mid-week messages are conditional to goal achievement: one-way dialogues are prompted if the user is on track to achieve their goal, a two-way dialogue with an open-ended question is prompted otherwise. To implement two-way dialogues, the scheduler uses Finite State Machines. Fig. 3 depicts the states of a dialogue for a goal accomplishment with 50 to 99%. Each state may comprise multiple messages. For instance, the initial state has five messages (i.e., greetings, restate current goal, goal progress summary, affirmation message, and goal modification) and one graphical plot (i.e., goal progress summary) in it.

Users transition to the subsequent state by choosing an option to either make the goal more challenging, less challenging, or maintaining the goal as is. Following the goal change, the users are prompted to indicate their confidence level in achieving the newly set goal. An indication of low confidence will prompt the user with potential barriers and strategies to achieve the agreed-upon goal. BeSMART holds three different FSMs which differ in the goal modification states. Goal modification states are conditional on PA goal achievement for the week.

4. Methods

The BeSMART coach aims to bolster participants' behavioral self-regulation by promoting self-monitoring of PA, establish progressive PA goals, and develop participant competence and support in achieving their PA goals. The BeSMART acceptability trial includes three key components: (1) Human health coaching, (2) Weekly SMS feedback, and (3) PA and blood pressure monitoring. The trial was conducted with $N = 10$ participants (9 female, 1 male) chosen randomly. To be eligible in the study, participants had to be mid-life adults, aged 35 to 64 years old, pre-hypertensive (120/80 mmHg–139/89 mmHg) or stage 1 hypertensive (140/90 mmHg–159/99 mmHg), and physically inactive with less than 100 min of MVPA the previous week. Participants were recruited from community and clinical feeder sites, including the university of Delaware's Office of Employee Health and Wellbeing, the University of Delaware's Nurse Managed Primary Care Center, and the Western Family YMCA. Recruitment strategies included flyers, social media, email, community events, word-of-mouth, and clinician referrals. All subjects were expected to have moderate CVD risk so medical clearance from their personal physician was required. The study was conducted for 4 weeks excluding the baseline and post-intervention weeks. This study was approved by the University of Delaware's Institutional Review Board (IRB 1771748-1) and all participants consented to participate in this 4 week study. Participants were asked to wear ActiGraphs to measure MVPA in baseline and post-intervention weeks. Participants measured their blood pressures everyday in the mornings for the duration of the entire study (6 weeks). During the intervention period, participants wore FitBits and were blinded to all FitBit notifications. Participants interacted with human health coaches twice during the study, once after the baseline week and once after the participant completed two weeks of intervention.

Table 1
Average Goal achievement, as a percentage of successfully completed bouts, across four weeks for each participant.

P101	P102	P103	P104	P105	P106	P107	P108	P109	P110
20.75	6.25	37.5	41.75	24.75	58.25	62.5	50	16.5	100

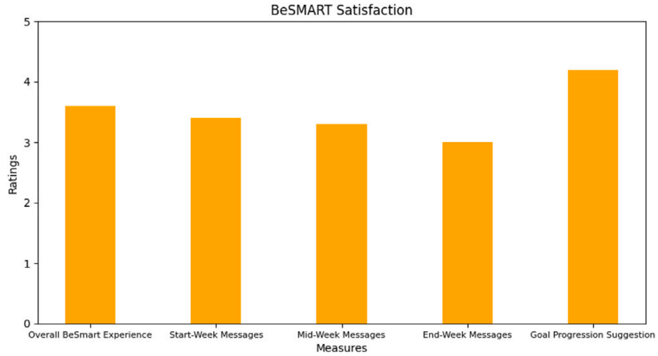

Fig. 4. Dialogue satisfaction.

5. Preliminary results

5.1. BeSMART goal achievement

To measure the effectiveness of our coaching system, we consider three outcomes: increase in moderate PA, change in blood pressure, and percentage of goals achieved during the coaching process. To see if the BeSMART system induced a behavior change, PA and blood pressure were compared from baseline and post intervention weeks. While no significant results were found, this was not expected of the acceptability trial. We also analyzed how participants performed during the coaching sessions.

Table 1 shows the average goal achievement across four weeks for each participant. This represents the percentage of total number of planned bouts successfully completed. A successful bout is defined as maintaining heart rate within the THR zone continuously for at least 80% of the duration of the planned activity. For instance, maintaining heart rate within THR for 16 or more consecutive minutes out of 20 min is considered a successful bout of 20 min. As goals change weekly, each user can have bouts of different lengths across weeks. The overall average of successful bouts across participants is 41.8% of planned bouts, and varied among participants from 6.25% to 100% with more than 60% of participants achieving less than 50% of their goals.

5.2. BeSMART text messaging satisfaction

To understand better the receptiveness of participants to BeSMART dialogues, we analyzed the exit survey data shown in Fig. 4. Fig. 4 depicts the satisfaction of participants on the scale of 1 to 5 with 5 being the highest satisfaction. The satisfaction with dialogues is categorized into satisfaction over: overall BeSMART experience (3.6), start-of-week (3.4), mid-week (3.3), end-of-week (3) and goal progression suggestions (4.2). The participants satisfaction can be considered good with the overall rating of 3.6. In the exit survey, when we asked specifically about the satisfaction of goal progression suggestions which are used in end-of-week messages, participants gave it a 4.2 score, liking how the BeSMART gave recommendations and participants were able to choose the goals they liked.

5.3. Transcript analysis from health coaching

To understand the type of preferences users have in PA domain, we extracted information pertaining to participants' preferences from health coaching session transcripts. 20 transcripts were generated from this study, 2 for each participant. Results from transcripts generated from the initial meetings with the health coaches were analyzed for participant preferences. The purpose of each session is to figure out participant's short/long term goals and action/coping plans to achieve these goals. To establish short/long term goals, health coaches ask participants to choose from an initial set of goals recommended by an exercise physiologist based on the participant's fitness level. Each user expressed a preference over what goal to choose by expressing preferences over duration, intensity and frequency of the goals. To achieve these goals, participants were asked to come up with activity and coping plans. Each participant expressed their preferences over activities, days, times, locations, and people who they do the activity with, to come up with action/coping plans. Table 2 summarizes the preferences of all the participants that we extracted. Based on the data, we categorize preference types in PA as follows:

Table 2
Preference Types in SMART PA domain.

Participant Details	Preference over Action plans						Preference over Alternatives			Preference of whom to do the activity with	Conditional Preferences
PID	No. of Activities	Activities	Location	Day/Time	Goal short term	Goal long term	Location	Day/Time	How?		Weather Work
101	1	Walking	Campus Workplace	Tuesday Thursday	60 min 2 days a week 30 min a day 3 bouts a day 10 min bouts 40%–44%		Neighborhood	Evenings	Motivate each other Walk in the building	Daughter	Work Companion Weather
102	1	Dance	Home	Tuesday, Friday	60 min 2 days a week 30 min a day 2 bouts a day 15 min bouts 40%–44%	150 min 5 days a week 30 min a day 1 bout a day 30 min bouts 45%–49%	Gym		Telling self that I need to do	Self	Time
103	1	Walking	Living room Neighborhood	Tuesday Saturday	40 min 2 days a week 20 min a day 2 bouts a day 10 min bouts 40%–44%			Next day Different time of day	Change location or activity	Dog or self	Weather
104	1	Walking	Basement	Monday Wednesday Saturday Probably mornings	66 min 3 days a week 22 min a day 1 bouts a day 22 min bouts 44%–49%	150 min 5 days a week 30 min a day 1 bout a day 30 min bouts 45%–49%			Just doing it Like the NIKE ad Setting up a reminder for particular days	Self	
105	1	Walking	Son's school On campus	Tuesday Thursday	60 min 2 days a week 30 min a day 3 bouts a day 10 min bouts 40%–44%	180 min 6 days a week 30 min a day 1 bout a day 30 min bouts 45%–49%		Next day	Reward self with particular podcast of interest Do treadmill at home if unable to make the time otherwise	Self	Weather Time
106	1	Walking	Work	Tuesdays Thursdays 10 am 4-5 pm	60 min 2 days a week 30 min a day 1 bouts a day 30 min bouts 40%–44%		Walmart Mall			Self	Weather
107	2	Walking Riding	Neighborhood Work	Mondays, Wednesdays Morning	60 min 2 days a week 30 min a day 2 bouts a day 15 min bouts 40%–44%	150 min 3 days a week 50 min a day 2 bout a day 25 min bouts 45%–49%				Self	Weather
108	1	Biking	Park/Trail	Tuesday, Friday, Saturday	75 min 3 days a week 25 min a day 1 bouts a day 25 min bouts 45%–49%	150 min 5 days a week 30 min a day 1 bout a day 30 min bouts 45%–49%	Stationary bikes in gym Dumbells at home	Wednesday		Self	Weather Work
109	1	Walking	Neighborhood	Monday, Wednesday	40 min 2 days a week 20 min a day 2 bouts a day 10 min bouts 45%–49%						
110	1	Walking	Neighborhood	Thursday, Friday, Sunday	75 min 3 days a week 25 min a day 1 bouts a day 25 min bouts 45%–49%	150 min 5 days a week 30 min a day 1 bout a day 30 min bouts 45%–49%			Double up another day	Self	Time Weather

5.3.1. *Preferences over activities*

Participants expressed absolute and conditional preferences over activities. Most of the participants apart from one preferred a single activity. But considering that our study is comprised by mid-life adults, when generalized, it is safe to assume people have preferences over multiple activity types. Participants also indicated these preferences were conditional on external factors like weather, location, time and other people. These can be generalized to people having preferences over different activities conditioned on different contexts. For instance, a person might walk when the weather is sunny, and when the location is home and run on a thread mill when the weather is rainy and the location is home.

5.3.2. *Preferences over short/long term goals*

Participants expressed preferences over short and long term goals. In particular, they expressed a preference over how their overall continuous numeric goal was discretized into smaller goals. In the PA domain, these preferences over short term goals evolve, eventually leading to a long term goal as PA behavior changes. For instance, participant 104's long term goal is 150 min per week decomposed into 30 min, 5 times a week. The short term goal is 66 min per week decomposed into 22 min, 3 times a week. In the case of successful behavior change, over a period of time, 66 min is evolved into 150 min, with emerging preferences of how each subgoal is decomposed.

Table 3
Average Goal achievement, as percentage of total planned time in THR, across four weeks for each participant.

P101	P102	P103	P104	P105	P106	P107	P108	P109	P110
59.58	100	100	78.38	64.58	100	100	100	85	100

Table 4
Total number of partial bouts (vs. complete bouts) in target heart rate.

P101	P102	P103	P104	P105	P106	P107	P108	P109	P110
3(2)	11(1)	23(5)	1(5)	7(2)	18(8)	36(8)	12(5)	6(1)	21(34)

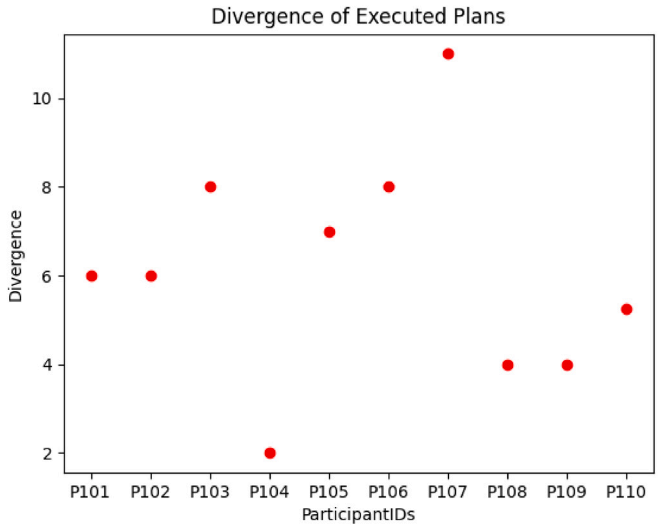

Fig. 5. Average weekly plan divergence.

5.3.3. *Preferences over people*

Only one participant expressed preferences for activities over other people. However, we generally assume that people would like to synchronize their PA with others. Essentially, these preferences imply that people coordinate their activities with others to achieve their goals.

6. Discussion

6.1. *Goal achievement as partial bouts*

To better understand the reasons behind low goal achievement as shown in Table 1, we examined the data by relaxing the notion of a successful bout. Table 3 shows the average goal achievement as a percentage of total minutes spent in moderate PA vs. total minutes planned that week. This is calculated by averaging the total minutes spent in THR ignoring pre-planned bouts. For instance, for a weekly goal of 3 bouts of 20 min at THR 100–115, we ignore the bouts and see the total minutes spent in THR (30 THR minutes out of 60 planned minutes translates to 50%). When this is calculated, the overall average goal achievement across participants is 81%, a 2x increase over the goal achievement as complete bouts. To see if these accumulated minutes spent in THR are unsuccessful attempts to complete goal bouts, we looked at partial bouts. A partial bout is defined as maintaining a heart rate within the THR zone continuously for more than 40% and less than 80% of the duration of the planned activity. The numbers in Table 4, show that the number of partial bouts are much greater than fully successful bouts (depicted in parenthesis). These results show that the difference in percentages in Table 4 (81%) and Table 1 (42%) is mostly because participants accomplished their goals in much smaller chunks than how health coaches recommended them to achieve these goals (fewer bouts and longer bouts). This suggests that executed action plans are much more dynamic and continuous in nature as opposed to initial action plans health coaches recommend.

6.2. *Plan divergence*

To quantitatively measure the difference between the initial and executed action plans, we found the divergence between the two. Divergence or plan difference (Fox, Gerevini, Long, & Serina, 2006) of initial and executed plans, is the number of actions that appear in the executed plan and not in the initial plan, plus the number of actions that appear in the initial plan and not in the executed plan. The actions in each plan correspond to the PA planned bouts or partial bouts. For instance, a weekly goal of 60 min of moderate PA in THR with 3 bouts of 20 min each has three actions where each action is composed of 20 min of PA in THR. Fig. 5 depicts the divergence measured of executed plans from initial plans for all participants. For each participant, divergence was measured across each week and averaged across four weeks. The overall average divergence across participants is 6 (rounded off to the nearest integer). The divergence varied from 2 to 11 across participants suggesting that: The initial action plans set at the beginning of the study are static and slight changes in execution violate these plans, but action plans should dynamically change and adapt to participant's current state to preserve the effects of having valid and specific action plans.

6.3. Dynamic action planing problem

One way to achieve such dynamism in action plans is for the coaching system to re-negotiate in the form of dialogues to figure out new plans. For instance, a user's action plan to achieve their PA goal of 60 min is walking 20 min three times a week on Mondays, Wednesdays and Fridays. Come Monday, the user might only partially fulfill their goal, walking 10 min in THR instead of 20. The system might hold a dialogue to ask the participant to come up with a new action plan to achieve the remaining 50 min. But to re-negotiate action plans with every change away from the initial plan, the automated coach system should conduct dialogues similar to the end-of-week dialogue multiple times through out the week. Given the average divergence of 6, there would be 6 extra negotiations on average, which might be a cause for concern. However, in the exit survey questionnaire shown in Fig. 4, participants *liked* goal progression suggestions giving them a score of 4.2. Participants liked how BeSMART gave recommendations and were able to choose the goals they liked. These recommendations by the system and choice given to user can bypass the dialogues but can still keep up with dynamic changes in action plans. For instance, in the example of 60 min divided into 3 bouts of 20 min. If the user only accomplishes 10 min, it would help the user better cope if the system gives options of either compensating this 10 min at another time on Monday itself, or some other day depending on the participant's preferred time. This automated computation of action plans dynamically will offload the burden of coming up with new action plans for every change in executed plans from the participants. To make such recommendations and choices, it is important that planning/scheduling systems take comprehensive preferences of users as shown in Table 2 into account while proposing new plans/schedules. To our knowledge, this problem has not been looked at in the literature. Below, we propose how to integrate these preferences into planning/scheduling frameworks to enable dynamic action planning.

7. Extensions to current work

The PA goals are numeric and continuous in nature. These goals can be further decomposed into smaller goals that can be achieved by users. To automate planning such goals, HTNs are best suited. In HTN planning (Ilghami, Munoz-Avila, Nau, & Aha, 2005; Nau et al., 2003), the planning system formulates a plan by decomposing tasks (symbolic representations of activities to be performed) into smaller and smaller subtasks until primitive tasks are reached that can be performed directly. An HTN planning problem consists of the following: the initial state (a symbolic representation of the state of the world at the time that the plan executor will begin executing its plan), the initial task network (a set of tasks to be performed, along with some constraints that must be satisfied), and a domain description that contains the following: a set of primitive actions, a set of methods that describe various possible ways of decomposing tasks into subtasks. A solution to a plan consists of a decomposed overall task into primitive actions with all the planning constraints satisfied. Though HTN plans are best suited to generate initial action plans, to adapt to changing plans of users, preferences need to be accommodated into the HTN generated plans. Preferences over alternative ways to achieve goals will help schedulers choose an alternative preferred action in case of total or partial execution failures. These preference extensions will allow schedulers to dynamically react to changing user contexts, avoiding the need to re-plan for every execution failure. Preferences extracted and categorized inSection 5 have allowed us to identify different extensions required to enable the dynamic execution of plans.

7.1. Methods with preferences

Extracted preferences from Table 4 indicate that the user's activity preferences are dependent on their context, particularly Location, Weather and Time. Given this knowledge, for real-time schedulers to pick the next most preferred activity, activities (methods) should have numeric preferences as one of the outcome criteria. Conditional preference networks such as UCP-nets (Boutilier et al., 2013) can serve as preference representations. The ability of UCP-nets to represent numeric conditional preferences along with linear optimization queries will allow schedulers to pick the next most preferred context for the activity in linear time. Fig. 6 represents a method, "Walk 30 min" with preference values that are conditionally dependent on weather, location, and time. In this example, sunny, home, and 6:00 am is the most preferred context to walk 30 min for a particular user. In case of execution failure, the scheduler will pick the next best context (sunny, work, 3:00 pm) as a recommendation to reschedule. These preference values might change for different activities, for different users.

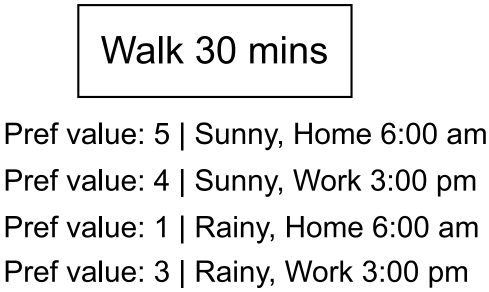

Fig. 6. Method with conditional preferences.

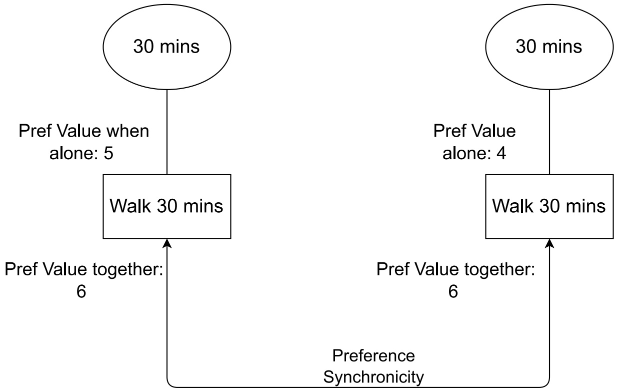

Fig. 7. Preference synchronization.

7.2. Tasks with preferences

In the PA domain, preferences over goals and how they are decomposed are prevalent. Each goal could be achieved in more than one way and users have a preference over how these goals are decomposed. Fig. 8 shows a 30 min PA goal with three different alternative decompositions, "Walk 15 min" twice, "Bike 30 min", and "Run 30 min". For schedulers to make a decision on which activities to choose from, it needs accumulation functions that consolidate preference values from the bottom (activities) to the root level (overall goal). We propose two such accumulation functions: PSum and PMax. PSum, when applied to a task, accumulates the preference values of all the activities present in the task. PMax, when applied to a task, selects the highest preference sub-task or activity present in the task. Fig. 8 shows PSum and PMax applied to a task structure with alternatives. The most preferred decomposition is chosen which is 30 min decomposed into two 15 min activities. The preference values of activities "Walk 15 min" are summed by PSum $5 + 5 = 10$, here the conditional preferences of activities are abstracted into a single value for simplicity. In case of execution failure of these activities, the next most preferred decomposition is chosen which is "Bike 30 min" for a recommendation. As accumulation functions assume the additive independence of preferences over activities, in more complicated cases these accumulation functions can also help switch between partially executed alternatives.

7.3. Coordination with preferences

It is not uncommon for people to plan their PA with others. Research indicates that people who worked in pairs or a group had higher success than those who did not (Hunter et al., 2015). For automated health coaches to assist users in scheduling activities with others, they need to coordinate with other automated health coaches. We propose three preferential coordination relationships: preferential synchronicity, hindrance, and facilitation. Fig. 7 shows an example of preference synchronicity over activities for two different individuals. The preference value of the activities increases to 6 from 5 and 4 when they are done simultaneously. These increases in preference values will encourage schedulers to coordinate their schedules with the schedules of other agents, thereby maximizing the overall preference values of schedules. Preferential facilitation and hindrance work in a similar way increasing or decreasing the preference values of activities when the facilitating/hindering activity is successfully executed.

7.4. Future work

One of the limitations of this study is limitations of the FitBit Charge 3 to detect activity types and monitor PA in real time. Knowing the type of activity will enable the automated coach to verify if the bouts accomplished by users are meant to be intended

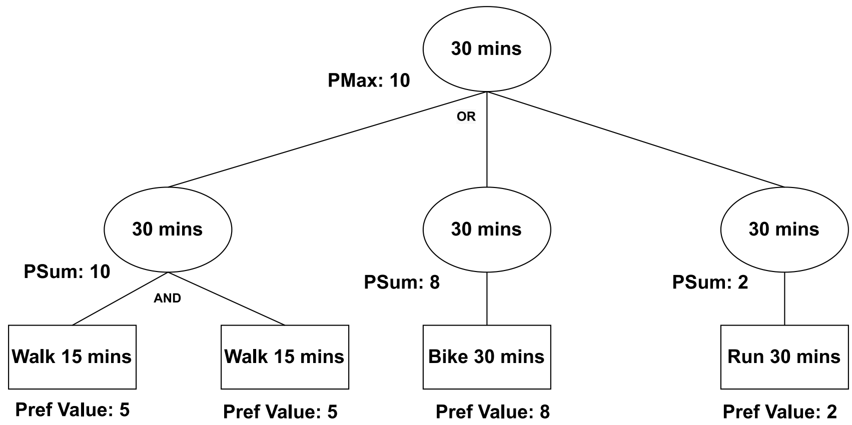

Fig. 8. Preference accumulation functions.

exercises. For future studies in which we plan to test our new planning/scheduling algorithms that solve the problem introduced in this paper, we intend to use Apple Watch platform which will enables us to detect activity types and also will allow us to provide Just-In-Time-Adaptive-Interventions (JITAI). Our other direction of research is to extend this problem to proactively plan/schedule to provide recommendations instead of re-actively. Proactive scheduling is quite useful in cases where the environment changes. For instance, it would be beneficial to proactively change action plans if the weather for the upcoming activity is not feasible to exercise. And lastly, to generate effective recommendations we need comprehensive preferences of users over various contexts and it might not be possible to elicit them verbally. We plan to learn these preferences by observing users.

8. Conclusion

In this paper, we introduced a novel semi-automated coaching system, BeSMART. We presented data from an acceptability study with 10 participants and identify a new problem which can improve the effectiveness of health coaching. From the experimental data, we identified that PA goals are numeric, continuous and their execution is quite dynamic requiring coaching systems to adapt to these changes in real time. In order to solve such a dynamic problem, we identified, extracted and categorized preferences that are common in PA and proposed extensions to current planning/scheduling frameworks which will enable new preference maximizing algorithms. The problem of dynamically generating action plans is not just applicable in the context of health coaching. Generating plans/schedules while maximizing user preferences can be applicable to all the domains where agents assist humans. For instance, in the Mars rover domain where plans/schedules to be executed by rovers are generally oversubscribed to resources, a subset of plans/schedules are chosen based on the preferences of scientists.

Declaration of competing interest

No author associated with this paper has disclosed any potential or pertinent conflicts which may be perceived to have impending conflict with this work. For full disclosure statements refer to https://doi.org/10.1016/j.smhl.2023.100389.

Data availability

The authors do not have permission to share data

Acknowledgments

We would like to thank Jennie Turner for helping us run this study. This study was funded by the National Institute on Aging, NIH R21AG056765-01.

References

Althoff, T., Clark, K., & Leskovec, J. (2016). Large-scale analysis of counseling conversations: An application of natural language processing to mental health. *Transactions of the Association for Computational Linguistics*, *4*, 463–476.

Bickmore, T. W., Schulman, D., & Sidner, C. (2013). Automated interventions for multiple health behaviors using conversational agents. *Patient Educ Couns.*, *92*(2), 142–148.

Bickmore, T. W., Trinh, H., Olafsson, S., O'Leary, T. K., Asadi, R., Rickles, N. M., et al. (2018). Patient and consumer safety risks when using conversational assistants for medical information: An observational study of Siri, Alexa, and Google Assistant. *Journal of Medical Internet Research*, *20*(9), Article e11510.

Boutilier, C., Bacchus, F., & Brafman, R. I. (2013). UCP-networks: A directed graphical representation of conditional utilities. arXiv preprint arXiv:1301.2259.

Boutilier, C., Brafman, R. I., Domshlak, C., Hoos, H. H., & Poole, D. (2004). CP-nets: A tool for representing and reasoning withconditional ceteris paribus preference statements. *Journal of Artificial Intelligence Research*, *21*, 135–191.

Brafman, R. I., & Chernyavsky, Y. (2005). Planning with goal preferences and constraints. In *ICAPS* (pp. 182–191).

Brafman, R. I., & Engel, Y. (2009). Directional decomposition of multiattribute utility functions. In *International conference on algorithmic decision theory* (pp. 192–202). Springer.

Decker, K. (1996). TAEMS: A framework for environment centered analysis & design of coordination mechanisms. *Foundations of Distributed Artificial Intelligence*, 429–448.

Dennis, S. M., Harris, M., Lloyd, J., Davies, G. P., Faruqi, N., & Zwar, N. (2013). Do people with existing chronic conditions benefit from telephone coaching? A rapid review. *Australian Health Review*, *37*(3), 381–388.

Düking, P., Tafler, M., Wallmann-Sperlich, B., Sperlich, B., Kleih, S., et al. (2020). Behavior change techniques in wrist-worn wearables to promote physical activity: Content analysis. *JMIR MHealth and UHealth*, *8*(11), Article e20820.

Eakin, E. G., Lawler, S. P., Vandelanotte, C., & Owen, N. (2007). Telephone interventions for physical activity and dietary behavior change: A systematic review. *American Journal of Preventive Medicine*, *32*(5), 419–434.

Fahrenwald, N. L., & Walker, S. N. (2003). Application of the transtheoretical model of behavior change to the physical activity behavior of WIC mothers. *Public Health Nursing*, *20*(4), 307–317.

Fox, M., Gerevini, A., Long, D., & Serina, I. (2006). Plan stability: Replanning versus plan repair. In *ICAPs, Vol. 6* (pp. 212–221).

Gonzales, C., & Perny, P. (2004). GAI networks for utility elicitation. *KR*, *4*, 224–234.

Gupta, I., Di Eugenio, B., Ziebart, B. D., Liu, B., Gerber, B. S., & Sharp, L. K. (2020). Goal summarization for human-human health coaching dialogues. In *FLAIRS conference* (pp. 317–322).

Gupta, I., Di Eugenio, B., Ziebart, B. D., Liu, B., Gerber, B. S., & Sharp, L. K. (2021). Summarizing behavioral change goals from SMS exchanges to support health coaches. In *Proceedings of the 22nd annual meeting of the special interest group on discourse and dialogue* (pp. 276–289).

Heintzelman, M. P., Dominick, G. M., Vemuri, A., & Decker, K. (2022). Development of the be smart feasibility trial to increase physical activity in midlife adults. In *Annals of behavioral medicine*: *vol. 56*, (no. SUPP 1), (p. S253). OXFORD UNIV PRESS INC JOURNALS DEPT, 2001 EVANS RD, CARY, NC 27513 USA.

Hunter, R. F., McAneney, H., Davis, M., Tully, M. A., Valente, T. W., & Kee, F. (2015). "Hidden" social networks in behavior change interventions. *American Journal of Public Health*, *105*(3), 513–516.

Ilghami, O., Munoz-Avila, H., Nau, D. S., & Aha, D. W. (2005). Learning approximate preconditions for methods in hierarchical plans. In *Proceedings of the 22nd international conference on machine learning* (pp. 337–344).

Jorge, A., McIlraith, S. A., et al. (2008). Planning with preferences. *AI Magazine*, *29*(4), 25.

Khatib, L., Morris, P., Morris, R., & Rossi, F. (2001). Temporal constraint reasoning with preferences.

Kinnafick, F.-E., Thøgersen-Ntoumani, C., & Duda, J. L. (2014). Physical activity adoption to adherence, lapse, and dropout: A self-determination theory perspective. *Qualitative Health Research*, *24*(5), 706–718.

Kivelä, K., Elo, S., Kyngäs, H., & Kääriäinen, M. (2014). The effects of health coaching on adult patients with chronic diseases: A systematic review. *Patient Education and Counseling*, *97*(2), 147–157.

Lentferink, A. J., Oldenhuis, H. K., de Groot, M., Polstra, L., Velthuijsen, H., & van Gemert-Pijnen, J. E. (2017). Key components in ehealth interventions combining self-tracking and persuasive ecoaching to promote a healthier lifestyle: A scoping review. *Journal of Medical Internet Research*, *19*(8), Article e277.

Lesser, V. R. (2002). Evolution of the GPGP/TAEMS domain-independent coordination framework. In *Proceedings of the first international joint conference on autonomous agents and multiagent systems: Part 1* (pp. 1–2).

Liao, P., Dempsey, W., Sarker, H., Hossain, S. M., Al'Absi, M., Klasnja, P., et al. (2018). Just-in-time but not too much: Determining treatment timing in mobile health. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, *2*(4), 1–21.

Mahon, S., Krishnamurthi, R., Vandal, A., Witt, E., Barker-Collo, S., Parmar, P., et al. (2018). Primary prevention of stroke and cardiovascular disease in the community (PREVENTS): Methodology of a health wellness coaching intervention to reduce stroke and cardiovascular disease risk, a randomized clinical trial.

Marcus, B. H., Bock, B. C., Pinto, B. M., Forsyth, L. A. H., Roberts, M. B., & Traficante, R. M. (1998). Efficacy of an individualized, motivationally-tailored physical activity intervention. *Annals of Behavioral Medicine*, *20*(3), 174–180.

Marcus, B. H., Rakowski, W., & Rossi, J. S. (1992). Assessing motivational readiness and decision making for exercise. *Health Psychology*, *11*(4), 257.

Márquez-Hernández, R., Hsu, L., McCoy, K., Decker, K., Vemuri, A., Dominick, G., et al. (2022). Towards development of an automated health coach. In *NLG4Health workshop*.

Miller, S. M., Brody, D. S., & Summerton, J. (1988). Styles of coping with threat: Implications for health. *Journal of Personality and Social Psychology*, *54*(1), 142.

Mitchell, E. G., Maimone, R., Cassells, A., Tobin, J. N., Davidson, P., Smaldone, A. M., et al. (2021). Automated vs. Human health coaching: Exploring participant and practitioner experiences. *Proceedings of the ACM on Human-Computer Interaction*, *5*(CSCW1), http://dx.doi.org/10.1145/3449173.

Muntaner, A., Vidal-Conti, J., & Palou, P. (2016). Increasing physical activity through mobile device interventions: A systematic review. *Health Informatics Journal*, *22*(3), 451–469. http://dx.doi.org/10.1177/1460458214567004.

Nau, D. S., Au, T.-C., Ilghami, O., Kuter, U., Murdock, J. W., Wu, D., et al. (2003). SHOP2: An HTN planning system. *Journal of Artificial Intelligence Research*, *20*, 379–404.

Oddone, E. Z., Gierisch, J. M., Sanders, L. L., Fagerlin, A., Sparks, J., McCant, F., et al. (2018). A coaching by telephone intervention on engaging patients to address modifiable cardiovascular risk factors: A randomized controlled trial. *Journal of General Internal Medicine*, *33*(9), 1487–1494.

op den Akker, H., Cabrita, M., op den Akker, R., Jones, V. M., & Hermens, H. J. (2015). Tailored motivational message generation: A model and practical framework for real-time physical activity coaching. *Journal of Biomedical Informatics*, *55*, 104–115. http://dx.doi.org/10.1016/j.jbi.2015.03.005.

op den Akker, H., Jones, V. M., & Hermens, H. J. (2014). Tailoring real-time physical activity coaching systems: A literature survey and model. *User Modeling and User-Adapted Interaction*, *24*(5), 351–392. http://dx.doi.org/10.1007/s11257-014-9146-y.

Peintner, B., & Pollack, M. E. (2004). Low-cost addition of preferences to DTPs and TCSPs. In *AAAI* (pp. 723–728).

Pérez-Rosas, V., Sun, X., Li, C., Wang, Y., Resnicow, K., & Mihalcea, R. (2018). Analyzing the quality of counseling conversations: The tell-tale signs of high-quality counseling. In *Proceedings of the eleventh international conference on language resources and evaluation*.

Reiter, E. (2020). Could NLG systems injure or even kill people? URL https://ehudreiter.com/2020/10/20/could-nlg-systems-injure-or-even-kill-people/.

Rossi, F., Van Beek, P., & Walsh, T. (2006). *Handbook of constraint programming*. Elsevier.

Rupp, M. A., Michaelis, J. R., McConnell, D. S., & Smither, J. A. (2018). The role of individual differences on perceptions of wearable fitness device trust, usability, and motivational impact. *Applied Ergonomics*, *70*, 77–87.

Ryan, R. M., Patrick, H., Deci, E. L., & Williams, G. C. (2008). Facilitating health behaviour change and its maintenance: Interventions based on self-determination theory. *European Health Psychologist*, *10*(1), 2–5.

Saponaro, M., Vemuri, A., Dominick, G., & Decker, K. (2021). Contextualization and individualization for just-in-time adaptive interventions to reduce sedentary behavior. In *Proceedings of the conference on health, inference, and learning* (pp. 246–256).

Shaparau, D., Pistore, M., & Traverso, P. (2006). Contingent planning with goal preferences. *vol. 21*, In *Proceedings of the national conference on artificial intelligence* (no. 1), (p. 927). Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999.

Smith, S. F., Gallagher, A., & Zimmerman, T. (2007). Distributed management of flexible times schedules. In *Proceedings of the 6th international joint conference on autonomous agents and multiagent systems* (pp. 1–8).

Sohrabi, S., Baier, J. A., & McIlraith, S. A. (2009). HTN planning with preferences. In *Twenty-first international joint conference on artificial intelligence*.

Spruijt-Metz, D., Wen, C., & O'Reilly, G. (2015). Innovations in the use of interactive technology to support weight management. *Current Obesity Report*, *4*(4), 510–519.

Svetkey, L. P., Batch, B. C., Lin, P.-H., Intille, S. S., Corsino, L., Tyson, C. C., et al. (2015). Cell phone intervention for you (CITY): A randomized, controlled trial of behavioral weight loss intervention for young adults using mobile technology. *Obesity*, *23*(11), 2133–2141.

Teixeira, P. J., Carraça, E. V., Markland, D., Silva, M. N., & Ryan, R. M. (2012). Exercise, physical activity, and self-determination theory: A systematic review. *International Journal of Behavioral Nutrition and Physical Activity*, *9*(1), 1–30.

Vemuri, A., Decker, K., Saponaro, M., & Dominick, G. (2021). Multi agent architecture for automated health coaching. *Journal of Medical Systems*, *45*(11), 1–7.

Wagner, T., & Lesser, V. (2000). Design-to-criteria scheduling: Real-time agent control. In *Workshop on infrastructure for scalable multi-agent systems at the international conference on autonomous agents* (pp. 128–143). Springer.

Wilson, P. M., Mack, D. E., & Grattan, K. P. (2008). Understanding motivation for exercise: A self-determination theory perspective. *Canadian Psychology/Psychologie Canadienne*, *49*(3), 250.

Wolever, R. Q., Simmons, L. A., Sforzo, G. A., Dill, D., Kaye, M., Bechard, E. M., et al. (2013). A systematic review of the literature on health and wellness coaching: Defining a key behavioral intervention in healthcare. *Global Advances in Health and Medicine*, *2*(4), 38–57.