

---

## Research and Applications

# An expandable approach for design and personalization of digital, just-in-time adaptive interventions

Suat Gonul,<sup>1,2</sup> Tuncay Namli,<sup>2</sup> Sasja Huisman,<sup>3</sup> Gokce Banu Laleci Erturkmen,<sup>2</sup> Ismail Hakki Toroslu,<sup>1</sup> and Ahmet Cosar<sup>1</sup>

<sup>1</sup>Department of Computer Engineering, Middle East Technical University, Ankara, Turkey, <sup>2</sup>SRDC Software Research & Development and Consultancy Corp., Ankara, Turkey, and <sup>3</sup>Department of Internal Medicine (Endocrinology), Leiden University Medical Center, Leiden, the Netherlands

Corresponding Author: Suat Gonul, PhD, Software Research & Development and Consultancy Corp., Silikon Blok Kat: 1 No: 16, ODTU Teknokent, 06800, Ankara, Turkey (suat@srdc.com.tr,suat.gonul@ceng.metu.edu.tr)

Received 26 March 2018; Revised 17 August 2018; Editorial Decision 2 November 2018; Accepted 15 November 2018

### ABSTRACT

**Objective:** We aim to deliver a framework with 2 main objectives: 1) facilitating the design of theory-driven, adaptive, digital interventions addressing chronic illnesses or health problems and 2) producing personalized intervention delivery strategies to support self-management by optimizing various intervention components tailored to people's individual needs, momentary contexts, and psychosocial variables.

**Materials and Methods:** We propose a template-based digital intervention design mechanism enabling the configuration of evidence-based, just-in-time, adaptive intervention components. The design mechanism incorporates a rule definition language enabling experts to specify triggering conditions for interventions based on momentary and historical contextual/personal data. The framework continuously monitors and processes personal data space and evaluates intervention-triggering conditions. We benefit from reinforcement learning methods to develop personalized intervention delivery strategies with respect to timing, frequency, and type (content) of interventions. To validate the personalization algorithm, we lay out a simulation testbed with 2 personas, differing in their various simulated real-life conditions.

**Results:** We evaluate the design mechanism by presenting example intervention definitions based on behavior change taxonomies and clinical guidelines. Furthermore, we provide intervention definitions for a real-world care program targeting diabetes patients. Finally, we validate the personalized delivery mechanism through a set of hypotheses, asserting certain ways of adaptation in the delivery strategy, according to the differences in simulation related to personal preferences, traits, and lifestyle patterns.

**Conclusion:** While the design mechanism is sufficiently expandable to meet the theoretical and clinical intervention design requirements, the personalization algorithm is capable of adapting intervention delivery strategies for simulated real-life conditions.

**Key words:** just-in-time adaptive interventions, personalized intervention delivery, reinforcement learning, digital intervention design, m-health

---

### OBJECTIVE

Adaptive interventions have emerged to deal with persons' varying responses in terms of adherence to and adoption of health behaviors that pertain to the treatment and self-management of chronic diseases and

health problems.<sup>1</sup> The technological progress in the past few years has enabled the delivery of adaptive interventions wherever and whenever they are needed via mobile devices. Interventions delivered in such a spontaneous way are known as just-in-time adaptive interventions (JITAs).<sup>2</sup>

Our objective is to deliver an expandable software framework for design and personalization of digital interventions in conformance with the JITAI framework.<sup>3</sup> Offering constructs matching with the JITAI components, namely, *tailoring variables*, *decision points*, *decision rules*, and *intervention options*, the proposed design mechanism facilitates JITAI design activities, as described in.<sup>4</sup> The mechanism is highly customizable and expandable with add-on constructs, enabling designers to customize the core capabilities to develop JITAI tailored to a particular health problem/population. The design constructs are bound to accompanying software modules, which can also be customized and expanded, for processing of heterogeneous data as desired.

The personalization module targets care receivers by monitoring their self-management and health-related behaviors and delivering mobile phone notifications in response to their continuously changing context during daily activities. Closely linked with the aforementioned JITAI components, we propose a novel learning method that tailors intervention delivery strategies dynamically in terms of intervention type, timing, and frequency using machine learning techniques, in compliance with people's action plans, changing physical/psychological contexts, as well as their changing preferences over time. We should note that these interventions do not aim to replace the care and support of healthcare professionals but aim to facilitate self-management following shared decision making, where care receivers and care providers agree on behavioral treatment goals and daily actions to pursue and follow-up until the next clinical visit of the care receiver.<sup>5</sup> We use the term "action plan" to refer to these behavioral goals and daily actions in the rest of the manuscript.

Concerning the validation, our objective is to show that the proposed design mechanism is sufficiently flexible and capable of addressing the requirements for JITAI design as specified in the literature. We also aim to show that the proposed personalization algorithm is able to adapt the intervention delivery according to individual needs. Specific to the personalization part, the goal is to validate the proposed approach in simulated settings before deploying it in a real-world case study.

## BACKGROUND AND SIGNIFICANCE

Approximately 70% of the deaths globally, corresponding to 40 million people, are due to chronic diseases.<sup>6</sup> Treatment of chronic conditions relies on people's self-management, which requires their ability to perform health-related behaviors, such as proper use of medication, physical activity, and following dietary recommendations, by themselves in the context of their daily lives.<sup>7</sup> Chronic diseases are a great burden on primary healthcare systems,<sup>8</sup> as average clinic visit times is only 10.3 minutes, leaving insufficient time for professional guidance to affected people for optimal self-management of their care.<sup>9</sup>

Emerging self-management programs have tried to fill the gap between self-management support in clinical visits and everyday life.<sup>10</sup> However, recent systematic reviews pointed out the lack of theoretical foundation of contemporary self-management support applications.<sup>11,12</sup> Furthermore, to be effective, self-management support programs must be tailored to suit people's priorities, resources, and lifestyles while taking multiple physiological and personal psychosocial factors into account.<sup>13</sup>

There are numerous recent studies of systems that provide personalized self-management support to people with chronic diseases<sup>14,15</sup> or health-related problems such as obesity<sup>16</sup> or sedentary behavior.<sup>17</sup> Waki et al.<sup>18</sup> present a system providing lifestyle

recommendations, matched to patients' inputs about food and exercise. The system gathers measurements twice a day. An intervention is determined and delivered right after data gathering. Gustafson et al.<sup>19</sup> present self-management modules, 1 of which tracks GPS data and warns people with (prior) alcohol addiction when they approach a previously identified high-risk location.

Such systems are usually static, rule-based systems evaluating intervention delivery conditions using the same rule set in every evaluation. Static systems are limited in terms of personalization of the intervention delivery strategies, as they do not adapt themselves in a systematic way to maintain engagement of people with interventions and extending adherence to a care program. Prolonged adherence is inversely proportional with the burden created by the interventions.<sup>20</sup> Due to their feedback loop-based learning mechanisms, dynamic systems can adapt intervention delivery according to personal values, conditions, or patterns, thus reducing the burden of interventions. In addition to the momentary contexts, some parameters might evolve throughout a long period of time. For example, a person's preferences and perceptions about the interventions might change over time. The receiver might get used to the interventions and form a kind of habituation towards similar interventions,<sup>21</sup> or might start feeling burdened as the number of interventions is too high and require too much cognitive resources.<sup>22</sup> Dynamic systems are able to recognize such changes without introducing additional rules.

There are research studies on computational approaches dealing with *adaptivity* and *just-in-timeness* of interventions, which are the 2 optimization dimensions that we also target.

Computational approaches deal with the *adaptivity* of interventions by recognizing longer-term changes in individuals. Some studies develop tailor-made models targeted at specific health behaviors and problems. For example, Chih et al.<sup>23</sup> present an agent-based model for uncovering the predictors of food choice and obesity in the presence of cue-reward conditioning. Similarly, Goldstein et al.<sup>24</sup> use supervised machine learning techniques to predict dietary lapses based on contextual tailoring variables and deliver personalized interventions according to the strong predictors of the lapse. In contrast, some dynamic system models lay out a more generic architecture capturing general, mathematic models of behavior change theories. Such systems are able to adapt JITAI parameters according to behavioral, model-based changes over time.<sup>25</sup> Navarro-Barrientos et al.<sup>26</sup> developed weight-loss interventions proposing a dynamic system model based on the Theory of Planned Behavior;<sup>27</sup> Martin et al.<sup>28</sup> propose a simulation of a model based on Social Cognitive Theory<sup>29</sup> to support physical activity of individuals over time.

Other studies combine mobile/sensing technologies with machine learning techniques to optimize intervention delivery with respect to timing of interventions. For example, Pejovic and Musolesi<sup>30</sup> utilize and compare a set of classifiers to predict opportune moments to deliver interventions based on specific variables including time, daily steps, location, and emotions. Boyer et al.<sup>31</sup> use Bayesian networks to predict drug cravings based on physiological parameters including skin conductance, skin temperature, motion, and pulse.

In<sup>32</sup> and,<sup>33</sup> authors utilize reinforcement learning for optimization of intervention delivery in real time by modeling it as a contextual bandit problem.<sup>34</sup> These studies are similar to ours in terms of representation of the real-time intervention delivery optimization problem with reinforcement learning constructs and validation of the proposed optimization algorithms in a simulated setting. Similar to our approach, the authors model the environment state with elements representing various momentary context parameters such as location, calendar data, or physical activity status. The authors present a simplified and

**Table 1.** Elements of JITAI design template

JITAI Component	Template Element	Description	Examples
-	<i>Targeted behavior</i>	The type of activity that the person is supposed to perform according to personal self-management goals and planned actions (ie, the action plan)	Blood glucose monitoring
-	<i>Description</i>	States the objective and reasoning of the JITAI along with the conditions suitable for delivering the intervention in a human-readable manner	One achieves a daily, weekly, or monthly blood glucose monitoring goal consecutively, and the system motivates her/him to maintain the behavior
<i>Decision points</i>	<i>Decision points</i>	Decision points for evaluating the decision rules of the associated JITAI. It can either be <u>event based</u> or <u>time bound</u> .	event = {upcoming_action, post_action} [ <i>Interpretation: Intervention options will be considered when a planned action has occurred (only motivations) or if there is an upcoming planned action (only reminders).</i> ]
<i>Intervention options</i>	<i>Behavior change technique (BCT)</i> <i>Content</i>	A reference to the behavior change technique introduced in the literature The message to be delivered to the person. It can include placeholders for injecting dynamically calculated information at the intervention delivery time. Placeholders may differ according to the BCT. Multilinguality is also supported.	Providing rewards contingent on successful behavior (derived from CALO-RE taxonomy <sup>38</sup> ) “en”: “ Well done you are doing a great job! You successively achieved your BG monitoring goal for last \${streak_value} \${streak_temporal}s.”, “es”: “ Bien hecho, está haciendo un gran trabajo! Su objetivo de monitorización de la glucosa ha sido alcanzado exitosamente durante los últimos \${streak_value} \${streak_temporal}.”, The example has two placeholders that are <i>streak_value</i> and <i>streak_temporal</i> . While the former placeholder represents the number of sequential temporal periods in which the person reached the blood glucose monitoring goal, the latter specifies the temporal period eg days, weeks or months.
<i>Decision rules / tailoring variables</i>	<i>Rule</i>	Decision rules that must be satisfied for delivering the intervention. This variable takes values conforming to the rule definition language described below.	[goal.monthly = ACHIEVED and goal.monthly[1]=ACHIEVED”, “goal.weekly = ACHIEVED and goal.weekly[1]=ACHIEVED”, “goal.daily = ACHIEVED and goal.daily[1]=ACHIEVED”]
<i>Distal / proximal outcomes</i>	<i>Associated goal</i>	Goals specify the targets to be achieved towards the ultimate clinical outcome. They are defined in action plans. Via this element the intervention instance is linked to 1 or more goals.	<ul style="list-style-type: none"> <li>• Monitoring blood glucose levels three times a day</li> <li>• Minimum 8000 steps per day</li> <li>• 7% HbA1c at the end of three months</li> </ul>

simulated real-life scenario on reducing smoking for heavy smokers as an instantiation of their conceptual model. The study considers a single generic intervention type and simulates decision points of JITAIs and people’s smoking urge. This approach differs from ours with respect to the learning algorithm. While we utilize Q-Learning<sup>35</sup> supported by transfer learning for cross-individual knowledge transfer,<sup>36</sup> the referred study utilizes an actor-critic algorithm.

van de Ven et al.<sup>37</sup> provide an expandable intervention design mechanism, with capabilities to define intervention decision rules bound to contextual parameters as well as time. The authors claim that, so far, their platform is the only one providing a flexible and versatile solution for definition and triggering interventions via mobile phones based on user behavior and context, which is consistent with our literature survey. However, that study does not provide any dynamic optimization on its intervention triggering mechanism.

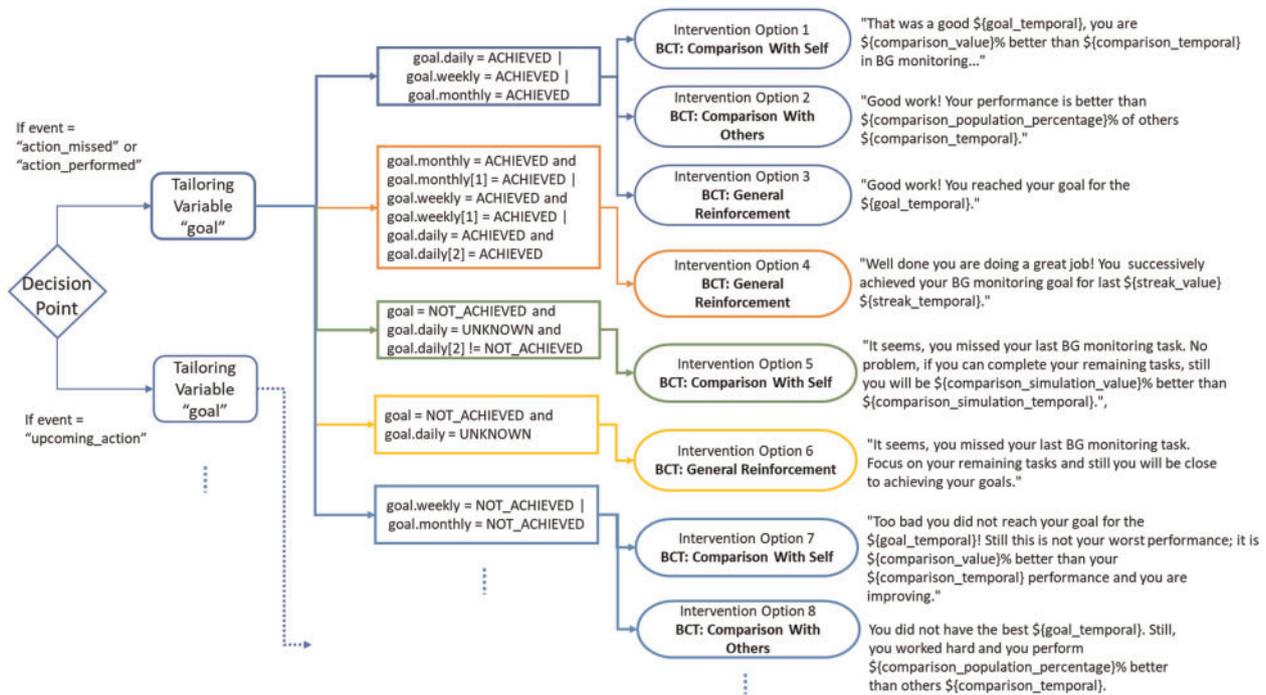
## MATERIALS AND METHODS

### Template-based JITAI design

As summarized in Table 1, we introduce a template-based design mechanism that enables configuration of interventions targeting

specific health problems in compliance with the JITAI framework presented in.<sup>3</sup> The template enables users to instantiate 4 JITAI components, namely, *decision points*, *intervention options*, *tailoring variables*, and *decision rules*.

The *decision point* element allows the definition of event-based and time-bound points. Time-bound points could be specified as specific points (eg, at 8:00 am and 9:00 pm) or periodic times (eg, at each 30 minutes between 10:00 am and 6:00 pm). Event-based points capture participant-initiated points, which could be manual (eg, when the person asks for care provider support) or bound to a change in a certain tailoring variable (eg, when the daily step count exceeds 10 000). Event-based points are also linked to the self-management goals and action plans of care receivers. Each planned activity in the action plan is supposed to be performed within a time frame (eg, “first blood glucose monitoring activity should be performed between 9:00 am and 11:00 am”). (See Section 2 of the [Supplementary Material](#) for a more detailed action plan). In this respect, the proposed framework proposes built-in event-based decision point implementations, namely, “*upcoming\_action*” and “*post-action*.” These concepts restrict the set of eligible interventions, such that only reminder interventions are triggered if there is



**Figure 1.** Example instantiation of the JITAI design constructs (ie, rule definition language elements). Overall, the figure shows the instantiation of rule definition language elements leading to several alternatives of motivation interventions. Decision points are the links connecting each intervention type to action plans. Considering the examples, all the intervention types are linked to the action plan slots classified as motivation.

an upcoming planned action; and only motivation interventions are triggered once a planned action has occurred. Figure 1 includes a more comprehensive example showing the relation between the action-plan-based decision points and the rest of the design constructs. The figure, overall, shows how the rule definition language elements can be used to define several alternatives of motivation JITAIs.

The *behavior change technique* and *content* elements enable specification of the *intervention options*. The *rule* element corresponds to the *decision rule* component. The *rule* element also incorporates the *tailoring variables* as operands in the decision rules. We elaborate on the rule construct in the next section, *Rule definition language*.

Nahum-Shani et al.<sup>3</sup> have introduced 2 additional concepts, namely, *distal outcome* and *proximal outcome* to represent goals to be achieved by care receivers with the support of JITAIs. Distal outcomes usually represent the ultimate goals as primary clinical outcomes such as losing weight or having lower levels of HbA1c. Proximal outcomes represent relatively short-term goals through which the effectiveness of interventions can be measured. These 2 concepts are not direct properties of a JITAI. Rather, they are captured by the goals defined in action plans (an example action plan is provided in Section 2 of the [Supplementary Material](#)). Each intervention instantiated via the proposed template is associated via the *associated goal* element. In this way, an intervention can be linked with one or more goals and, therefore, proximal and distal outcomes. As described in the *Rule definition language* section, we provide built-in functionalities for measuring the effectiveness of interventions considering the targeted outcomes.

**Rule definition language**

A multidisciplinary discussion between computer scientists and cognitive behavioral psychologists concluded that the interventions must meet the following requirements to be effective. They should

1) ensure clinical safety, capturing the specifications from clinical guidelines, 2) be in line with personal self-management goals and actions as planned in clinical visits, 3) be in line with health behavior change theories, and 4) conform to people’s preferences in order not to create burden with irrelevant notifications sent at inappropriate times.<sup>51</sup> Aiming to be sufficiently expandable towards meeting these requirements, we propose a rule definition language with the following Backus–Naur form notation:<sup>39</sup>

```

<tailoring_variable> ::= <tailoring_variable> |
  <tailoring_variable><temporal>
<temporal> ::= <temporal><temporal> " [" <index> "]"
<rule> ::= <tailoring_variable><operator><number> |
  <tailoring_variable><operator><tailoring_variable>
<rule_list> ::= <rule>j<rule><boolean_op><rule_list>
<rule_list_list> ::= <rule_list>j<rule_list> ", "
  <rule_list_list>
<decision_rule> ::= " [" <rule_list_list> "]"
    
```

**Listing 1 grammar of the rule definition language**

*Contexts*, representing tailoring variables, are data integrating and processing constructs, with corresponding software modules. They transform data aggregated from external sources to the required format. Contexts also perform information extraction operations to produce the actionable information as a scalar and actionable value by the rule. For example, we introduce a context named *adherence* for quantification of how close the person is to achievement of the self-management goal associated with the behavior. Goal-related contexts allow gauging the effectiveness of interventions. However, each goal type has its own internal logic for such an evaluation, which requires specialized methods. Contexts can be specialized by suffixing sub-contexts. For example, *adherence: bgm* outputs the adherence for the blood glucose

monitoring goal, as long as the corresponding software module is in place.

The *temporal* construct is used to evaluate a context considering a specific time interval. It can be set to *daily*, *weekly*, or *monthly* to get the average values during the specified interval. It can also be set to *best*, *weekly-worst* kind of specifiers to get peak values for the desired context. The temporal construct can also be assigned with an index that allows data retrieval for a specific period in the past. As an example, regarding the design constructs introduced, the design language supports the following expression:  $stress.monthly < stress.monthly [-1]$ . The expression is interpreted as follows: The average stress value in the current month should be less than the average stress value during the last month.

The proposed approach is expandable in terms of integration of additional data sources, either static or streaming, eg, demographic information, self-reported information via questionnaires, measurements obtained from medical devices, or contextual information sensed from wearables. It is also expandable with additional contexts, which could be used to implement even more specific data processing or calculation operations. These data management modules can then be reused to design new JITAIs.

Section 1 of the [Supplementary Material](#) contains the complete list of built-in constructs included in the rule definition language, as specifically used in the POWER2DM study/real word care program, for which we present the details later.

### Personalization of JITAIs

The goal of the JITAI personalization mechanism is to increase the effectiveness of the interventions with respect to their timing (*just-in-time-ness*) and frequency and type (*adaptivity*). Benefiting from machine learning techniques, the proposed personalization method puts the proposed approach beyond rule-based systems, as it is able to adapt itself according to continuously changing contexts and personal variables including both long-term (eg, past performance, habit strength, preferences, etc.) and short-term parameters (eg, location, time, etc.).

### Analogy between reinforcement learning and personalization of JITAIs

We use reinforcement learning<sup>40</sup> for personalization of the intervention delivery, as the conceptual elements of RL and the problem of interest match perfectly as depicted with an analogy in [Figure 2](#). A person and associated context data correspond to the environment entity. The context information at a certain time represents the state of the environment. The changing context of a person in daily life is modeled as a series of state transitions of the environment. Interventions delivered to the person correspond to the actions performed by the learning agent. The engagement of the person with interventions is emitted as a reward signal reflecting the acceptance of the delivered intervention by the person.

tions delivered to the person correspond to the actions performed by the learning agent. The engagement of the person with interventions is emitted as a reward signal reflecting the acceptance of the delivered intervention by the person.

### Advancing the analogy with the JITAI components

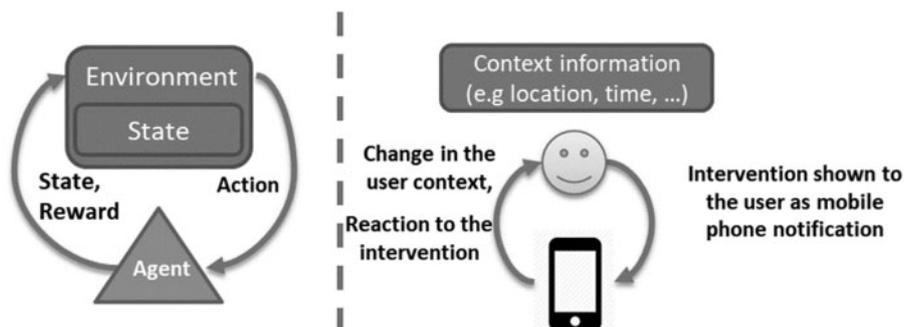
Besides the design mechanism, the reinforcement learning models can also be aligned with the JITAI components, letting them work in harmony with the literature-driven, conceptual foundations of JITAIs. We claim that the reinforcement learning approach can capture the dynamics of these elements only if 2 dedicated reinforcement learning models are employed simultaneously. In this respect, we propose 2 models named *intervention-selection* and *opportune-moment-identification* aiming at optimization of type (and frequency) and timing of JITAIs respectively.

Both the intervention-selection and the opportune-moment-identification steps have their own *decision points*. Considering a person's daily action plan, each scheduled activity is a decision point for intervention-selection. For the opportune-moment-identification case, we treat each moment when there occurs a change in the person's momentary context as a decision point.

*Intervention options* such as type, content, or timing are characteristics of a JITAI. These characteristics are distributed over the 2 learning models. The intervention-selection model optimizes the type and frequency of JITAIs via the selected action at each state change, whereas the opportune-moment-identification model captures the timing of JITAIs.

Similarly, *tailoring variables* are distributed over the learning models. Tailoring variables form the state element of the reinforcement learning models. While the states of the intervention-selection model are composed mostly of the long-term parameters (eg, intervention engagement history or intervention preferences), the states of the opportune-moment-identification model include mostly the short-term parameters (eg, step count or location).

*Decision rules* are constructed with conditional expressions with tailoring variables as operands, as seen in the example in [Table 1](#). On the reinforcement learning side, the states contain all the tailoring variables. Therefore, state-action mappings accumulated inside the learning agent's policy are a collection of decision rules. Each mapping in the policy includes the action to be taken in a state. This means that each distinct state leads to a decision based on the conditions represented with the instantiations of tailoring variables. The agent optimizes its policy over time by learning from people's experiences. As it visits different states, it learns how to behave in different conditions and finds the (near-)optimal strategy that suits the



**Figure 2.** Analogy between a traditional RL setup and intervention delivery optimization problem. While the left part shows the elements of an RL setup along with the information flow between them, the right part includes the corresponding elements and information flow concerning the optimization of intervention delivery.

user best. With this base approach, though, the learning algorithm needs to consider all tailoring variables inside the person state against the complete set of actions (ie, intervention types), which would require a long learning time. Instead, decision rules associated with individual intervention instances are used to limit the complete action set by keeping only the interventions of which rules are satisfied considering the momentary contextual values. For example, The JITAI defined in Table 1 is eligible only if the person reached his/her goal 3 times consecutively in 1 of the specified time frames.

The amount of rewards collected by the learning agent during the learning process is an indicator for achieving the goals set for the behavior targeted by the interventions. As we presented earlier, the goals are representatives of the proximal/distal outcomes desired. An indirect relation between rewards and proximal/distal outcomes can be established, such that the more rewards that the learning agent collects, the more the targeted outcomes are achieved.

### Overall JITAI personalization algorithm

The JITAI personalization algorithm implements the conceptual approach described throughout the analogy above. It first identifies the eligible interventions by evaluating their decision rules. At the second step, it selects the best-matching intervention (or selects none of the eligible interventions) based on the current context and historical experiences of the person. After the intervention is selected, its content is finalized by calculating dynamic variables included in the content. The algorithm then discovers an opportune moment to deliver the intervention. The process goes on with the update of the learning models. Figure 3 elaborates on the algorithm in detail.

### Simulation testbed

To validate the personalization of intervention delivery strategies, we lay out a simulation testbed for simulating the treatment and self-management process of people with diabetes. The testbed is composed of 3 main concepts to be simulated, namely, an *action plan*, *JITAI*s, and *personas*. We perform care process simulations for 2 personas with differentiating conditions related to the simulated concepts. Our aim is to show that the JITAI personalization algorithm captures the persona-specific conditions and adapts the intervention delivery accordingly. We present the details about the simulated concepts below.

#### Action plan

The simulated *action plan* is similar for both personas. It includes blood glucose monitoring activity 3 times a day, before meals, within pre-set timeframes. The link between action plans and intervention types, as described below, is established via their *targeted behavior*. That is, only the interventions targeting the same behavior with the action plan, among all available interventions, are considered for delivery to support the activity included in the action plan. See Section 2 of the [Supplementary Material](#) for more details about the simulated action plan.

#### JITAI

Targeting the blood glucose monitoring behavior, we have 3 JITAI instances, each of which implements a specific behavioral change technique (BCT) as described in the CALO-RE taxonomy. Below, we present distinguishing characteristics of their components. Fur-

thermore, Section 3 of the [Supplementary Material](#) shows their instantiation via the proposed design mechanism.

- *Intervention-1: Prompting self-monitoring of behavior*: The first intervention is a standard reminder. Having the reminder category, the intervention is supposed to be sent within the period during which the activity is supposed to be performed.
- *Intervention-2: Reminding with comparing with others*: This is also a reminder intervention. While reminding of the activity, it also motivates the person by presenting a comparison with others in terms of performance of the targeted behavior. This intervention has the same decision rules as the previous one. So, both will be considered as eligible interventions at the same decision points.
- *Intervention-3: Praising the performed behavior*: This one is a motivational intervention complimenting the person on successful performance of the planned behavior. It is associated with 3 decision rules representing the achievement of the monitored goal in daily, weekly, and monthly timescales, respectively. Meeting only 1 of them is sufficient for making this intervention eligible for delivery.

#### Personas

Personas themselves differ further in 3 dimensions: 1) *commitment intensity*, as the motivation felt by a person to perform a behavior, 2) *preferences on intervention types*, and 3) *daily activities*. Each dimension is elaborated below:

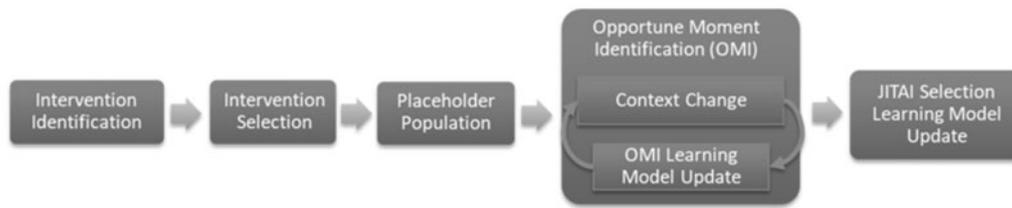
*Social-Psychological Model of Prospective Memory and Habit*: To test the JITAI personalization algorithm with respect to the adaptive frequency of interventions throughout the care process, we utilize the concept of habit formation.<sup>40</sup> Though many self-monitoring behaviors cannot be seen as habitual behaviors strictly, we made use of the empirically validated, mathematical model of habit formation<sup>41</sup> to construct a mathematical expression of the correlation between commitment and actual behavior. It is assumed that, during the care and evolving self-management process, people become less dependent of extrinsic reminders and motivators. When people respond to the delivered interventions and the behavioral goal is obtained consistently, they will need fewer reminding notifications, so interventions need to become more intermittent. To calculate and simulate parameters concerning the performance of a *new behavior*, we apply the mathematical model of habit formation to simulate care processes and calculate individual *habit strength* scores, which express the likelihood of performing the *new behavior*, without being dependent on extrinsic reminders and motivators. The model requires an initial commitment intensity value for each person for its internal calculations. Therefore, we use commitment intensity as 1 of the differentiating factors of the simulated personas. Table 2 presents the commitment intensities for each persona along with possible interpretations.

#### Preferences on intervention types

We associate personas with preferences on the 3 intervention types, as presented in Table 3, as a reflection of personal acceptance on them. Each value represents the probability of engaging with the specified intervention by the person when one is encountered.

#### Simulation of daily activities

Characteristics of daily activities also affect people's reactions to interventions. We simulate daily activities by selecting type, occurrence, and duration of activities randomly from pre-defined daily ac-



1. *Identify the eligible interventions for delivery.* (This step is triggered for each activity prescribed in the action plan of the person.)
  - a. Retrieve the scheduled activity from action plan considering the time of the day (e.g. blood glucose monitoring).
  - b. Identify the complete set of interventions associated with the retrieved behavior. (For example, a motivational intervention can be instantiated in different ways; e.g. by providing information about the activities performed, comparison-by-self or comparison-with-others[38].)
  - c. Evaluate the decision rules of the identified interventions.
  - d. Mark interventions of which decision rules are evaluated positively, as eligible.
2. *Select Intervention*
  - a. Compose the intervention-selection model state with the adaptivity-related parameters e.g. habit strength, behavior frequency, obtained from the current person state
  - b. Select the intervention by evaluating the previous experiences related to the current state. (In this step not delivering any intervention is also an option. So, this means that the algorithm would also decide on delivering a intervention or not in this step)
    - i. If the current state was encountered before, select the intervention with the highest score
    - ii. If the current state was not encountered before,
      1. Identify the similar states concerning the state parameters
      2. Calculate the average rewards for each intervention throughout the similar states
      3. Select the intervention with the highest reward
3. Calculate the *placeholder* values, if any is included in the intervention content
4. *Identify the best opportune moment* for intervention delivery, if a intervention was selected in the previous step. (This step is triggered upon a change on the person state until the intervention is delivered within the time frame associated with the current activity.)
  - a. Compose the opportune-moment-identification model state with the just-in-timeness-related parameters e.g. physiological, psychological or environmental contextual parameters including time of the day, physical activity status, etc. obtained from the current person state.
  - b. Decide to deliver the intervention or not by evaluating the previous experiences related to the current state
    - i. If the current state was encountered before, select the intervention with the highest score
    - ii. If the current state was not encountered before,
      1. Identify the similar states concerning the state parameters
      2. Calculate the average rewards for each intervention throughout the similar states
      3. Select the intervention with the highest reward
  - c. Update the aggregated reward for the state – selected action pair of the opportune-moment-identification model where the action could be delivering or not delivering the intervention according to the reaction to intervention and performance of the behavior.
5. Update the aggregated reward for the state – selected action pair of the intervention-selection model where the selected action could a certain type of intervention or no intervention according to the reaction to intervention and performance of the behavior.

**Figure 3.** Overall JITAI personalization algorithm. The flow at the top of the figure shows the main steps of the algorithm executed sequentially. First, the set of eligible interventions is identified; then the algorithm selects 1 of the eligible interventions considering current context and past experiences. The placeholders are populated, if there are any. Next is the identification of the best moment to deliver the intervention. Finally, the learning models are updated based on persons' engagement with interventions.

tivity timeline templates. A template includes a set of activities placed on a timeline. For each simulated day, a new set of activities is generated, such that each activity is represented with a set of parameters as exemplified in Table 4. Intervention engagement rules

in the simulation are established, such that engagement occurs only if the person is physically sedentary, mentally neutral, and the phone screen is on. Simulating the performance of the targeted behavior differs from simulation of engagement with the intervention. We

**Table 2.** Persona's commitment intensities and their interpretations

	Commitment Intensity [0, 1.0]	Possible Interpretation of the Commitment Intensity
Person-1	0.7	An indicator of giving more importance to the targeted behavior considering the expected eventual benefits
Person-2	0.3	An indicator that s/he takes the behavior change less seriously and giving less importance to the behavior as s/he does not expect eventual benefit from performing the behavior; or s/he perceives the behavior as a relatively complex / challenging task. The same task might be perceived in varying difficulties by different people <sup>42,43</sup>

**Table 3.** Preference indicators for JITAIs

	Intervention-1 (reminder 1)	Intervention-2 (reminder 2)	Intervention-3 (motivation)
Person-1	50%	70%	70%
Person-2	50%	0%	10%

**Table 4.** Parameters of a simulated daily activity

activity_description	Takes a break sitting on the couch and watching TV
location	HOME
physical_activity	SEDENTARY
phone_usage	ACTIVE
phone_check_suitability	Yes
state_of_mind/emotional_status	(One of the following values is chosen randomly according to the associated probability) CALM/RELAXED=30% CALM/NEUTRAL=20% FOCUS/RELAXED=25% FOCUS/NEUTRAL=20% TENSE/NEUTRAL=5%
start_time	Relative to previous activity
start_time_variation	0
duration	45 minutes
duration_variation	15 minutes (ie, the duration might change between 30 and 60 minutes)
behavior_performance_suitability	Yes

simply mark certain activities in the timeline as suitable for behavior performance by setting the “*behavior\_performance\_suitability*” as also seen in the table. Please see Section 7 of the [Supplementary Material](#) to see an example sequence of simulated daily activities.

Driven by the benchmarked parameters of simulated personas, we introduce the following hypotheses as conditions for which the JITAI personalization algorithm is expected to adapt intervention delivery accordingly. For each hypothesis, we present corresponding simulation parameter leading to the hypothesis:

- *Varying parameter:* According to the habit formation model used, people with higher commitment intensities perform the

behavior more frequently and reach maximum habit strength faster.

*Hypothesis 1:* Person-2 should receive interventions for a longer time than Person-1, as performing the new behavior habitually will take longer for Person-2.

- *Varying parameter:* A well-formed habit (ie, automatic performance of the *new behavior*) indicates that the person performs the behavior with less dependence on extrinsic reminders and motivators.

*Hypothesis 2:* Intervention delivery frequency should decrease throughout the simulated care process.

- *Varying parameter:* As a reflection of their individual differences and preferences, the simulated people favor different intervention types.

*Hypothesis 3:* The number of delivered intervention for each intervention type should be proportional to people's preferences for them.

- *Varying parameter:* Daily activities of each person are generated semi-randomly based on personal activity timeline templates. Therefore, in addition to the distinct activities of each person, the activities for the same person vary among the simulated days because of the randomness included in the activity generation mechanism. Varying daily activities determine both when the behavior could be performed and when the intervention (ie, the mobile phone notification) can be engaged with.

*Hypothesis 4-a:* Interventions should be delivered in moments that are suitable for intervention engagement as specified by the simulation configurations.

*Hypothesis 4-b:* The intervention delivery should be temporally close to the periods that are marked suitable for behavior performance.

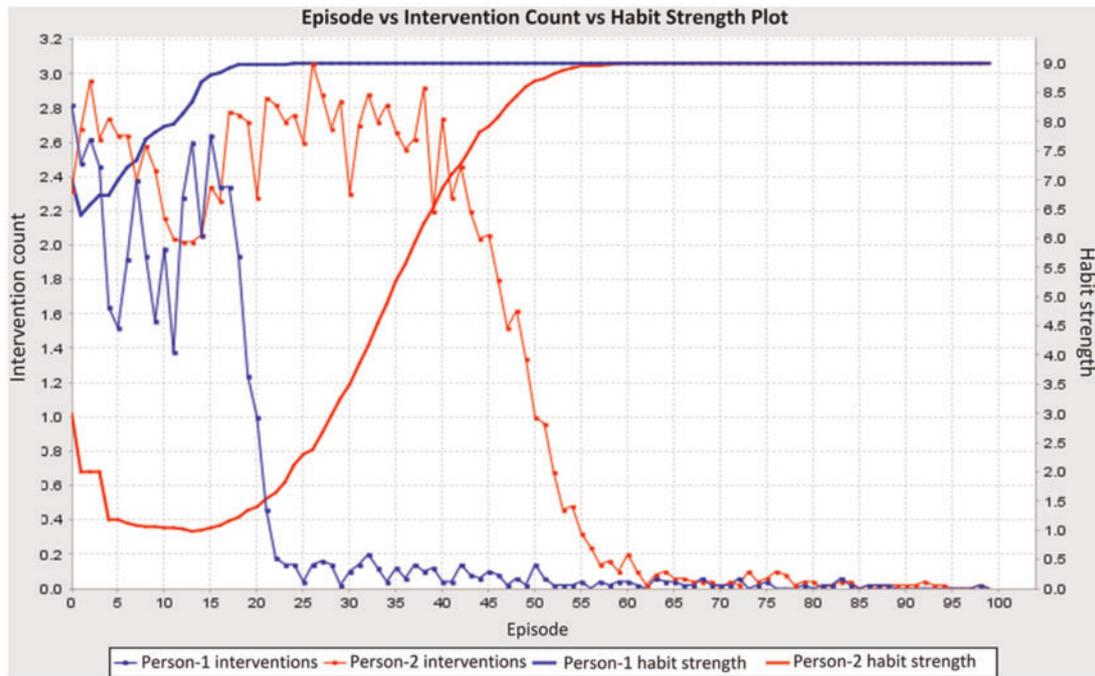
## RESULTS

We first show how the expandable JITAI design mechanism meets the JITAI design-related specifications derived from various resources such as BCT taxonomies, clinical guidelines, or algorithms for automated self-management support. Then, we run the simulation testbed, present the results, and discuss the results with respect to the proposed hypotheses.

### Validating the JITAI design mechanism

Clinical guidelines such as the American Diabetes Association (ADA)<sup>44</sup> guidelines for prevention or delay of Type 2 diabetes, the Joslin Clinical Guideline for Adults with Diabetes,<sup>45</sup> BCT taxonomies, eg, CALO-RE,<sup>38</sup> and algorithms related to self-management support, eg, Predictive 303,<sup>46</sup> provide clear starting points for JITAI components, even though the relation between the JITAI framework and those resources is not stated explicitly. For example, ADA guidelines recommend interrupting prolonged sitting every 30 minutes with short bouts of physical activity. In this example, the tailoring variable is the physical activity status, the decision point is every minute, and the decision rule would be to check whether the person has been inactive for the last 30 minutes.

A JITAI-focused analysis of these resources revealed that automating evidence-based self-management support strategies first requires integration of data for decision-making processes related



**Figure 4.** Episode vs. intervention count/habit strength plot. This plot shows the inversely proportional relation between the habit strength and number of interventions delivered.

to the personalization of JITAIs. The data vary in nature (structured/unstructured, static/streaming) and in source (devices, sensors, mobile app, questionnaires). To enable the analysis of the integrated data, they must be transformed into a format as required by decision rules. Concerning the interrupting inactivity case above, an example transformation would be reducing the last 30 minutes of physical activity data to a binary result. Finally, the transformed information should be injected into the decision-making processes of self-management strategies.

As validation of the design mechanism, we claim that the proposed approach meets exactly such JITAI-design-related requirements, as the core design constructs can be expanded with additional, reusable data integration and processing modules and decision rules for JITAI design. We have also validated the JITAI design capabilities in the POWER2DM (Predictive Model-Based Decision Support for Diabetes Patient Empowerment) Project,<sup>1</sup> aiming to develop a personalized self-management support system for diabetes patients. In POWER2DM, computer scientists and cognitive behavioral psychologists as the authors of this study have collaboratively developed interventions to support the self-management of diabetes patients regarding blood glucose monitoring, exercise, medication adherence, and carbohydrate monitoring. The [Supplementary Material](#), from Section 4 to Section 6, presents JITAI components from both theoretical and case-driven perspectives.

### Simulation results

We simulate 100 trials for each persona in which each trial lasts 100 days. Below, we elaborate how the obtained results validate the hypotheses above.

**Hypothesis-1 validation:** Figure 4 is a compact figure validating both Hypothesis-1 and Hypothesis-2 at the same time. The x-axis

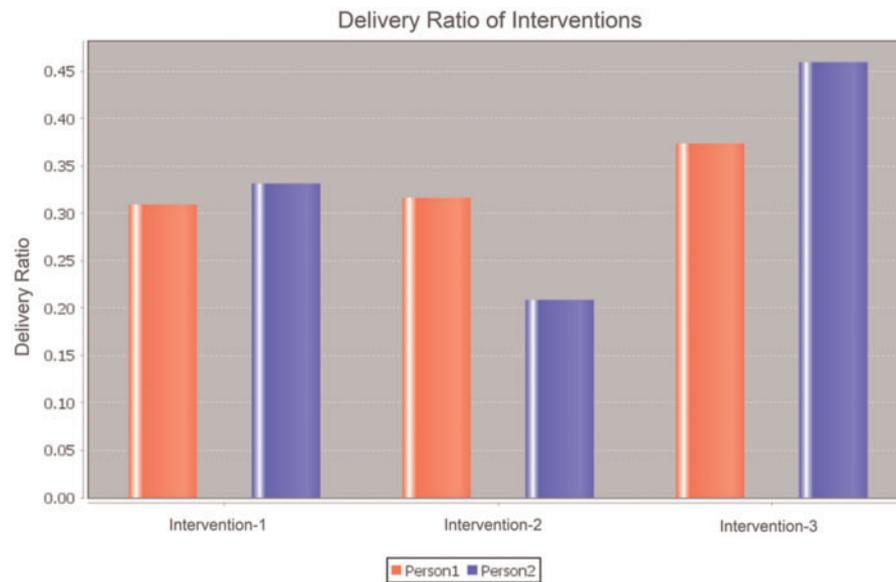
represents the simulated days in chronological order. The shaped lines show the simulated habit strength values as generated by the habit formation model at each simulated day. The blue and red lines show values for Person-1 and Person-2, respectively. On the other hand, the plain blue and red lines show the average number of interventions for Person-1 and Person-2. Accordingly, the left and right y-axes represent the value ranges for the average number of delivered interventions and simulated habit strength.

As configured with a relatively higher commitment intensity than Person-2, Person-1 reaches to maximum habit strength earlier, indicating that Person-1 forms a habit for the blood glucose monitoring behavior faster. The duration of intervention delivery is proportional to the length of the period through which the persons reach the maximum habit strength. These outcomes are also consistent with the results obtained in,<sup>47</sup> in which the authors develop a habit formation model on empirical data that outputs the habit formation duration for real-world behaviors with varying complexities. That study reveals that habit formation takes approximately 20 days for relatively less complex behaviors, but it increases to a few months for more complex behaviors. Person-2 can be seen as one who perceives the task as relatively complex, as we assumed in the simulation specifications above.

**Hypothesis-2 validation:** The average number of delivered interventions (shaped lines) is relatively high in the initial episodes. Despite the fluctuations, the number of interventions decreases as the simulated habit strength increases throughout the trials. The algorithm is also sensitive to the change in habit strength, such that the faster the habit strength reaches its maximum value, the faster the number of interventions delivered decreases.

**Hypothesis-3 validation:** Figure 5 shows the ratio of the number of interventions for each intervention type to the total number of interventions delivered for each intervention type for each persona. As Person-1 is neutral for Intervention-1 and slightly favors the other 2, the ratios presented in the figure are almost equal, but still

1 <http://www.power2dm.eu/>



**Figure 5.** Person vs. intervention type ratio plot. The plot shows the ratio of the number of a specific intervention type to the total number of interventions delivered for each intervention type for each persona.

reflect Person-1's preferences. On the other hand, although Person-2 is neutral only for Intervention-1 and has little interest for Intervention-3, Intervention-3 has the highest delivery ratio. The results for Intervention-1 and Intervention-2 do reflect the preferences of the person on reminder interventions.

The ratio of Intervention-3 is relatively high for 2 reasons: First, Intervention-3 instances are delivered in states in which it is predicted that the person would not remember to perform the behavior. This is expected, as we encouraged intervention delivery as a reminder cue regardless of the intervention type, in cases when it is likely the person would not remember to perform the behavior. (The results for Person-1 do not present such a pattern, as there are almost no states in which Person-1 forgets performing the behavior.) The second reason for the high ratio of Intervention-3 is that the algorithm is not able to learn the person preferences at the beginning of the learning phase and selects interventions randomly.

**Hypothesis-4 validation:** Figure 6 shows the difference between intervention delivery and behavior performance times. Almost 60% of the reminder interventions are delivered at most 30 minutes before performing the behavior. In general, as the time difference gets higher, the ratio of delivered interventions decreases except the "61-120 minute" bar. The bulge of this bar arises mainly from the high frequency of the daily activities that are convenient for intervention delivery at the corresponding time period during the day. For example, Person-1 has a working session, lasting 90 minutes on average, during which s/he regularly checks the phone. This actually means that the algorithm learns and exploits such common patterns in people's daily lives.

The results comply with the simulated intervention engagement rules related to the contextual parameters, such that 83% of interventions were sent when persons had a convenient emotional status including neutral, relaxed, or happy; 74% of interventions were sent when persons were in sedentary mode, and 63% of interventions were sent when the phone screen was on.

## DISCUSSION

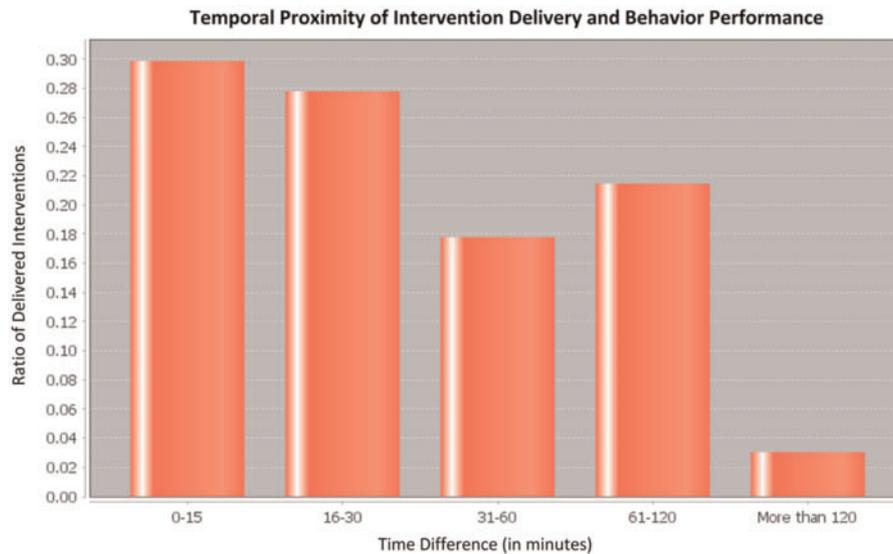
**Legitimacy of the approach:** The dynamic JITAI personalization mechanism and multi-dimensional expandability of the JITAI design

mechanism are the 2 innovative characteristics of the proposed approach, advancing the state-of-the-art research. Both characteristics bring opportunities for adoption of the system by various entities such as mobile/web application vendors in the behavioral health software market, public health organizations, or other healthcare organizations working on clinical studies aiming large-scale digital interventions at patient populations.

**Capturing the rules associated with the simulated concepts:** We already discussed the simulation results in the previous section by describing how they are aligned with the hypotheses and how they deviate from the expected results along with the causes of deviations. As a summary, we claim that the proposed algorithm is able to capture the rules that are associated with the simulated concepts.

We also presented the core algorithm that is used to optimize the intervention delivery policy in,<sup>48</sup> where we aimed to break the sedentary behaviors of office workers during working hours. In that study, we obtained better results for a machine-learning-based personalization mechanism compared to results obtained for interventions delivered according to a fixed schedule. The results were collected from both real-world and simulated test cases.

**Further improvements:** Despite the innovative character of the current approach, there is room for potential improvements. For example, as we mentioned earlier, the fluctuations of the intervention counts in Figure 4 happen when the learning algorithm encounters unknown states, ie, the cold-start problem. Tackling the cold-start problem, the evidence from various micro-randomized trials,<sup>49</sup> measuring the effect of individual intervention components, or any expert-knowledge-based heuristic can be utilized to provide a warm start for the learning algorithm. An unknown state might be encountered, eg, when the person attains a certain habit strength, the highest one, for the first time. Instead of taking random actions in such cases, the algorithm might employ a machine learning classifier to make an educated guess, or it might simply favor not delivering an intervention in proportion with the current habit strength. The reinforcement learning methodology is convenient for integration of such external knowledge by setting the initial scores of relevant state-action pairs inside the learning agent's policy accordingly. Furthermore, the learning models can be enriched with additional



**Figure 6.** Difference between intervention delivery and behavior performance times. Each bar represents the amount of time difference and the ratio of interventions delivered in that frame to the total interventions.

state parameters, eg, performance parameters to represent the person context more accurately.

Besides improvements to the learning models, the JITAI design capabilities could be improved with additional intervention types or content presentation modalities. Section 1 of the [Supplementary Material](#) presents all the currently available constructs of the design approach. Although a simple system with the limited number of built-in constructs targeting the POWER2DM case study, the design mechanism lays out the basis to expand the system with more constructs as needed by the targeted health problem.

Finally, in this study, we consider preferences for intervention types, commitment intensities, and daily activities as differentiating factors in the persona simulation. A more realistic simulation could be achieved by also considering factors such as self-efficacy, motivation, and prior experience, changing the behavior or outcome expectancies.

## CONCLUSION

We present a framework for JITAI design and personalization that can be customized for care programs targeting varying health problems and populations. The design mechanism, incorporating a rule definition language, can be specialized with add-on constructs to conceive interventions addressing the specific requirements of a care program. The personalization part employs a reinforcement learning-based approach to optimize/personalize the intervention delivery concerning the frequency, type, and timing of interventions dynamically according to data aggregated for a person over time.

We validate the JITAI design mechanism by providing example JITAI definitions in which the characteristics of JITAIs are extracted from various relevant resources available in the literature, such as clinical guidelines and taxonomies of behavior change, and by a real-world case study providing self-management support to diabetes patients.

We validate the personalized intervention delivery mechanism through a simulation testbed in which action plans, JITAIs, and personas, with differentiating characteristics, are simulated. We present that the personalization algorithm is able to capture the rules

associated with the simulated concepts, indicating its potential to be used in real-world settings. In future studies, we aim to validate the personalization mechanism empirically throughout a randomized controlled trial to be carried out in the scope of POWER2DM with 280 diabetes patients in total.<sup>50</sup>

## FUNDING

The research leading to these results has received funding partially from the European Community's H2020 Programme under grant agreement no. H2020-PHC-689444, POWER2DM project (Predictive Model-Based Decision Support for Diabetes Patient Empowerment) and partially from The Scientific and Technological Research Council of Turkey.

## CONTRIBUTORS

SG wrote the manuscript. SG, TN, and GBL are the leading authors formulating the conception of the work. SG developed the reinforcement learning-based personalization algorithms. SG, TN, and SH worked on the requirement analysis and definition of JITAIs. TN formulated and implemented the decision rule language. SG prepared the simulation setup and generated results. In general, the study was conducted under the supervision of GBL, IHT, and AC; they provided substantial feedback to the manuscript.

## SUPPLEMENTARY MATERIAL

[Supplementary material](#) is available at *Journal of the American Medical Informatics Association* online.

*Conflict of interest statement.* The authors have no competing interests to declare.

## REFERENCES

1. Almirall D, Nahum-Shani I, Sherwood NE, *et al.* Introduction to SMART designs for the development of adaptive interventions: with application to weight loss research. *Transl Behav Med* 2014; 4 (3): 260–74.

2. Spruijt-Metz D, Nilsen W. Dynamic models of behavior for just-in-time adaptive interventions. *IEEE Pervasive Comput* 2014; 13 (3): 13–7.
3. Nahum-Shani I, Smith SN, Spring BJ, et al. Just-in-time adaptive interventions (JITAI) in mobile health: key components and design principles for ongoing health behavior support. *Ann Behav Med* 2018; 52 (6): 446–62.
4. Nahum-Shani I, Hekler EB, Spruijt-Metz D. Building health behavior models to guide the development of just-in-time adaptive interventions: a pragmatic framework. *Health Psychol* 2015; 34(5): 1209–19.
5. Elwyn G, Frosch D, Thomson R, et al. Shared decision making: a model for clinical practice. *J Gen Intern Med* 2012; 27 (10): 1361–7.
6. Noncommunicable Diseases. World Health Organization 2017. <http://www.who.int/mediacentre/factsheets/fs355/en/> Accessed August 16, 2018.
7. Roberts CK, Barnard RJ. Effects of exercise and diet on chronic disease. *J Appl Physiol* 2005; 98 (1): 3–30.
8. Quinn CC, Gruber-Baldini AL, Shardell M, et al. Mobile diabetes intervention study: testing a personalized treatment/behavioral communication intervention for blood glucose control. *Contemp Clin Trials* 2009; 30 (4): 334–46.
9. Iyengar V, Wolf A, Brown A, et al. Challenges in diabetes care: can digital health help address them? *Clin Diabetes* 2016; 34 (3): 133–41.
10. Whitehead L, Seaton P. The effectiveness of self-management mobile phone and tablet apps in long-term condition management: a systematic review. *J Med Internet Res* 2016; 18 (5): e97.
11. Holmen H, Wahl AK, Cvancarova Smástuen M, et al. Tailored communication within mobile apps for diabetes self-management: a systematic review. *J Med Internet Res* 2017; 19 (6): e227.
12. Hood M, Wilson R, Corsica J, et al. What do we know about mobile applications for diabetes self-management? A review of reviews. *J Behav Med* 2016; 39 (6): 981–94.
13. Funnell MM, Anderson RM. Empowerment and self-management of diabetes. *Clin Diabetes* 2004; 22 (3): 123–7.
14. Ben-Zeev D, Brenner CJ, Begale M, et al. Feasibility, acceptability, and preliminary efficacy of a smartphone intervention for schizophrenia. *Schizophr Bull* 2014; 40 (6): 1244–53.
15. Fioravanti A, Fico G, Salvi D, et al. Automatic messaging for improving patients engagement in diabetes management: an exploratory study. *Med Biol Eng Comput* 2015; 53 (12): 1285–94.
16. Laing BY, Mangione CM, Tseng CH, et al. Effectiveness of a smartphone application for weight loss compared to usual care in overweight primary care patients. *Ann Intern Med* 2014; 161 (Suppl 10): S5–12.
17. Van Dantzig S, Geleijnse G, van Halteren AT. Toward a persuasive mobile application to reduce sedentary behavior. *Pers Ubiquit Comput* 2013; 17 (6): 1237–46.
18. Waki K, Fujita H, Uchimura Y, et al. DialBetics: a novel smartphone-based self-management support system for type 2 diabetes patients. *J Diabetes Sci Technol* 2014; 8 (2): 209–15.
19. Gustafson DH, McTavish FM, Chih M-Y, et al. A smartphone application to support recovery from alcoholism: a randomized clinical trial. *JAMA Psychiatry* 2014; 71 (5): 566–72.
20. Pellegrini CA, Pfammatter AF, Conroy DE, Spring B. Smartphone applications to support weight loss: current perspectives. *Adv Health Care Technol* 2015; 1: 13–22.
21. Pop-Eleches C, Thirumurthy H, Habyarimana JP, et al. Mobile phone technologies improve adherence to antiretroviral treatment in a resource-limited setting: a randomized controlled trial of text message reminders. *AIDS* 2011; 25 (6): 825–34.
22. Chen IJ, Chang CC. Content presentation modes in mobile language listening tasks: English proficiency as a moderator. *Comput Assist Lang Learn* 2011; 24 (5): 451–70.
23. Chih M-Y, Patton T, McTavish FM, et al. Predictive modeling of addiction lapses in a mobile health application. *J Subst Abuse Treat* 2014; 46 (1): 29–35.
24. Goldstein SP, Evans BC, Flack D, et al. Return of the JITAI: applying a just-in-time adaptive intervention framework to the development of m-health solutions for addictive behaviors. *Int J Behav Med* 2017; 24 (5): 673–82.
25. Riley WT, Cesar A, Rivera DE. The importance of behavior theory in control system modeling of physical activity sensor data. *Conf Proc IEEE Eng Med Biol Soc* 2014; 2014: 6880–3.
26. Navarro-Barrientos JE, Rivera DE, Collins LM. A dynamical model for describing behavioural interventions for weight loss and body composition change. *Math Comput Model Dyn Syst* 2011; 17 (2): 183–203.
27. Ajzen, I. (1985). From intentions to actions: A theory of planned behavior. In: J. Kuhl, & J. Beckman (Eds.), *Action control: From cognition to behavior* (pp. 11–39). Heidelberg: Springer.
28. Martin CA, Rivera DE, Riley WT, et al. A dynamical systems model of social cognitive theory. *Proc Am Control Conf* 2014; 1:2407–12.
29. Bandura A. *Social Foundations of Thought & Action: A Social Cognitive Theory*. Englewood Cliffs, NJ: Prentice-Hall, Inc.; 1986.
30. Pejovic V, Musolesi M. InterruptMe: designing intelligent prompting mechanisms for pervasive applications. In: AJ Brush, Adrian Friday, Julie Kientz, James Scott, Juneha Song (Eds.), *Proceedings ACM International Conference Ubiquitous Computing*; Seattle, WA, USA: ACM New York, NY, USA; 2014: 897–908.
31. Boyer EW, Fletcher R, Fay RJ, et al. Preliminary efforts directed toward the detection of craving of illicit substances: the iHeal project. *J Med Toxicol* 2012; 8 (1): 5–9.
32. Lei H, Tewari A, Murphy S. An actor-critic contextual bandit algorithm for personalized interventions using mobile devices. In: Max Welling, Zoubin Ghahramani, Corinna Cortes, Neil D. Lawrence, Kilian Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems*; Montreal, Canada: Neural Information Processing Systems (NIPS); 2014: 27.
33. Lei H, Tewari A, Murphy SA. An actor-critic contextual bandit algorithm for personalized mobile health interventions. *arXiv preprint* 2017; arXiv:1706.09090.
34. Langford J, Zhang T. The epoch-greedy algorithm for multi-armed bandits with side information In: Daphne Koller, Dale Schuurmans, Yoshua Bengio, L'On Bottou (Eds.), *Advances in Neural Information Processing Systems*; Vancouver, British Columbia, Canada: Neural Information Processing Systems (NIPS); 2008: 817–24.
35. Watkins CJ, Dayan P. Q-learning. *Mach Learn* 1992; 8 (3–4): 279–92.
36. Lazaric A. Transfer in reinforcement learning: a framework and a survey. In: Marco Wiering, Martijn van Otterlo (Eds.), *Reinforcement Learning*. Berlin, Heidelberg: Springer; 2012:143–73.
37. van de Ven P, O'Brien H, Henriques R, et al. ULTEMAT: a mobile framework for smart ecological momentary assessments and interventions. *Internet Interv* 2017; 9: 74–81.
38. Michie S, Ashford S, Sniehotta FF, et al. A refined taxonomy of behaviour change techniques to help people change their physical activity and healthy eating behaviours: the CALO-RE taxonomy. *Psychol Health* 2011; 26 (11): 1479–98.
39. Backus JW. The syntax and semantics of the proposed international algebraic language of the Zurich ACM-GAMM conference. In: S. de Picciotto (Editor), *Proceedings of the International Conference on Information Processing*; Paris, France: UNESCO (Paris); 1959: 125–32.
40. Sutton RS, Barto AG. *Reinforcement Learning: An Introduction*. Cambridge: MIT Press; 1998.
41. Tobias R. Changing behavior by memory aids: a social psychological model of prospective memory and habit development tested with dynamic field data. *Psychol Rev* 2009; 116 (2): 408–38.
42. Hom HL, Maxwell FR. The impact of task difficulty expectations on intrinsic motivation. *Motiv Emot* 1983; 7 (1): 19–24.
43. Scasserra D. *The Influence of Perceived Task Difficulty on Task Performance* [theses and dissertations]. 2008: 756. <http://rdw.rowan.edu/etd/756> Accessed August 16, 2018.
44. American Diabetes Association. Prevention or delay of type 2 diabetes. *Diabetes Care* 2017; 40 (Suppl 1): S44–7.
45. Joslin Clinical Guideline for Adults with Diabetes. Joslin diabetes center 2017. <https://www.joslin.org/docs/CLINICAL-GUIDELINE-FOR-ADULTS-WITH-DIABETES-Rev-05-17-2017.pdf> Accessed August 16, 2018.
46. Meneghini L, Koenen C, Weng W, et al. The usage of a simplified self-titration dosing guideline (303 Algorithm) for insulin detemir in patients

- with type 2 diabetes—results of the randomized, controlled PREDICTIVE™ 303 study. *Diabetes Obes Metab* 2007; 9 (6): 902–13.
47. Lally P, Van Jaarsveld CH, Potts HW, *et al.* How are habits formed: modelling habit formation in the real world. *Eur J Soc Psychol* 2010; 40 (6): 998–1009.
  48. Gonul S, Namli T, Baskaya M, *et al.* Optimization of Just-in-Time Adaptive Interventions Using Reinforcement Learning. In: Malek Mouhoub, Samira Sadaoui, Otmane Ayt Mohamed, Moonis Ali (Eds.), *The 31st International Conference on Industrial, Engineering & Other Applications of Applied Intelligent Systems (IEA-AIE 2018)*; Montreal, Canada: Springer; 2018.
  49. Klasnja P, Hekler EB, Shiffman S, *et al.* Microrandomized trials: an experimental design for developing just-in-time adaptive interventions. *Health Psychol* 2015; 34S: 1220–8.
  50. Sont J, Smith I, Uitbeijerse B, *et al.* POWER2DM—evaluation campaign methodology [project deliverable]. 2017. <http://www.power2dm.eu/wp-content/uploads/Power2DM-D5.3.pdf> Accessed August 16, 2018.
  51. Plößnig M, Smith I, Huisman S, *et al.* POWER2DM—dynamic behaviour change intervention models for self-management [project deliverable]. 2016. <http://www.power2dm.eu/wp-content/uploads/Power2DM-D3.1-1.pdf> Accessed October 29, 2018.