



UvA-DARE (Digital Academic Repository)

Reinforcement Learning to Send Reminders at Right Moments in Smartphone Exercise Application

A Feasibility Study

Wang, S.; Sporrel, K.; van Hoof, H.; Simons, M.; de Boer, R.D.D.; Ettema, D.; Nibbeling, N.; Deutekom, M.; Kröse, B.

DOI

[10.3390/ijerph18116059](https://doi.org/10.3390/ijerph18116059)

Publication date

2021

Document Version

Final published version

Published in

International Journal of Environmental Research and Public Health

License

CC BY

[Link to publication](#)

Citation for published version (APA):

Wang, S., Sporrel, K., van Hoof, H., Simons, M., de Boer, R. D. D., Ettema, D., Nibbeling, N., Deutekom, M., & Kröse, B. (2021). Reinforcement Learning to Send Reminders at Right Moments in Smartphone Exercise Application: A Feasibility Study. *International Journal of Environmental Research and Public Health*, 18(11), [6059].
<https://doi.org/10.3390/ijerph18116059>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)



Article

Reinforcement Learning to Send Reminders at Right Moments in Smartphone Exercise Application: A Feasibility Study

Shihan Wang ^{1,2,*}, Karlijn Sporrel ³ , Herke van Hoof ¹, Monique Simons ⁴, Rémi D. D. de Boer ⁵, Dick Ettema ³, Nicky Nibbeling ⁶ , Marije Deutekom ⁷ and Ben Kröse ^{1,5}

- ¹ Informatics Institute, University of Amsterdam, 1090 GH Amsterdam, The Netherlands; h.c.vanhoof@uva.nl (H.v.H.); b.j.a.krose@hva.nl (B.K.)
 - ² Department of Information and Computing Sciences, Utrecht University, 3584 CC Utrecht, The Netherlands
 - ³ Department of Human Geography and Spatial Planning, Utrecht University, 3584 CS Utrecht, The Netherlands; k.sporrel@uu.nl (K.S.); d.f.ettema@uu.nl (D.E.)
 - ⁴ Consumption & Healthy Lifestyles Group, Wageningen University & Research, 6700 HB Wageningen, The Netherlands; monique.simons@wur.nl
 - ⁵ Digital Life Centre, Amsterdam University of Applied Science, 1091 GC Amsterdam, The Netherlands; r.d.d.de.boer@hva.nl
 - ⁶ Centre of Expertise Urban Vitality, Amsterdam University of Applied Science, 1097 DZ Amsterdam, The Netherlands; n.nibbeling@hva.nl
 - ⁷ Faculty of Health, Sports and Welfare, Inholland University of Applied Sciences, 2015 CE Haarlem, The Netherlands; marije.deutekom@inholland.nl
- * Correspondence: s.wang2@uu.nl



Citation: Wang, S.; Sporrel, K.; van Hoof, H.; Simons, M.; de Boer, R.D.D.; Ettema, D.; Nibbeling, N.; Deutekom, M.; Kröse, B. Reinforcement Learning to Send Reminders at Right Moments in Smartphone Exercise Application: A Feasibility Study. *Int. J. Environ. Res. Public Health* **2021**, *18*, 6059. <https://doi.org/10.3390/ijerph18116059>

Academic Editor: Florian Fischer

Received: 6 April 2021

Accepted: 31 May 2021

Published: 4 June 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Abstract: Just-in-time adaptive intervention (JITAI) has gained attention recently and previous studies have indicated that it is an effective strategy in the field of mobile healthcare intervention. Identifying the right moment for the intervention is a crucial component. In this paper the reinforcement learning (RL) technique has been used in a smartphone exercise application to promote physical activity. This RL model determines the ‘right’ time to deliver a restricted number of notifications adaptively, with respect to users’ temporary context information (i.e., time and calendar). A four-week trial study was conducted to examine the feasibility of our model with real target users. JITAI reminders were sent by the RL model in the fourth week of the intervention, while the participants could only access the app’s other functionalities during the first 3 weeks. Eleven target users registered for this study, and the data from 7 participants using the application for 4 weeks and receiving the intervening reminders were analyzed. Not only were the reaction behaviors of users after receiving the reminders analyzed from the application data, but the user experience with the reminders was also explored in a questionnaire and exit interviews. The results show that 83.3% reminders sent at adaptive moments were able to elicit user reaction within 50 min, and 66.7% of physical activities in the intervention week were performed within 5 h of the delivery of a reminder. Our findings indicated the usability of the RL model, while the timing of the moments to deliver reminders can be further improved based on lessons learned.

Keywords: reinforcement learning; just-in-time adaptive intervention; reminder; physical activity; mobile application

1. Introduction

Inactive lifestyle and a lack of physical activity can lead to serious physical and mental health issues [1]. People are therefore advised to engage in regular physical activity (i.e., at least 150 min of moderate-intensity activity every week for adults [2]). However, many individuals struggle to maintain a healthy activity level. Consequently, the promotion of healthy physical activity behavior is still an open challenge [3]. Interventions using mobile exercise applications are considered promising for supporting physical activity behaviors,

as mobile devices are well integrated into people's daily lives [4] and can continuously deliver interventions to users [5].

Recently, just-in-time adaptive intervention (JITAI) has gained attention and been considered to be an effective strategy in the field of mobile healthcare intervention [6]. JITAI is an intervention design aiming to provide the right type and amount of support, at the right time, by adapting to an individual's changing internal and external state [7]. In the JITAI framework, identifying the right moment for intervention is a crucial component. As argued by Fogg, in daily life, the personal state that the user is involved in when an intervention is delivered can play an important role in determining its effectiveness [8]. For instance, a reminder about physical activity sent during a meeting will likely be ineffective, no matter how persuasive its content might be.

JITAI in mobile health applications are mainly used to prevent certain health threats, including addictive behaviors such as overeating [9,10], smoking [11] and prolonged sedentary behaviors [12]. In other words, that research mainly concentrates on finding 'vulnerable states' in which a user is urged not to perform a certain behavior. A few recent systems have started to research the right moments of 'opportunity' to promote a certain behavior, although they are still at the early stage. Ding et al. [13] described the potential of such a system and demonstrated how sending reminders at the right moments could improve the persuasiveness of mobile interventions. In their study, researchers set predetermined opportunistic moments in which the user is invited to walk (e.g., when the individual is overusing a smart phone or when he or she is sedentary). Expanding on this previous research, we proposed the incorporation of a data-driven 'moment of opportunity' in a smartphone application to promote user physical activity.

However, identifying the 'opportunity' moments from data to deliver interventions in mobile applications is demanding, as the right moments for an individual can change according to personal momentary needs. One possible solution is to use the reinforcement learning (RL) technique, as it can continuously adapt the intervention strategy based on both the momentary state and the feedback of users [14]. RL has been used in mobile applications to support a certain physical activity. Most researchers have continuously adapted the *content or type* of interventions by tracking the momentary behaviors of people. For instance, Rabbi et al. [5] recommended different types of physical activity adapting to users' changing needs using a bandit algorithm. Yom-Tov et al. [15] personalized the contents of messages for patients with diabetes to encourage their physical activity. Zhou et al. [16] delivered interventions to users suggesting adaptive and personalized daily step goals. In an intervention focused on weight loss, Forman et al. [17] optimized the combination of text messages and human coaching intervention based on a cumulative reward score. Conversely, several researchers have concentrated on using RL to adapt to the *moments* of intervention delivery with respect to users' on-going status and behavior. In one recent work, Liao et al. [18] developed the HeartSteps application to actively decide whether to provide interventions at five fixed moments per day. However, too many interactions over a short period (e.g., 5 times per day) could add burden for the user and adversely impact engagement [19]. Driven by this practical issue, we allow a limited number of reminders within each week in our smartphone application, and use RL to identify the 'right' moments for sending a restricted number of reminders. As far as we know, this has never been studied in mobile health intervention settings.

In this paper, we concentrate on motivating users to perform a physical activity by delivering reminders at personalized moments. The moments when reminders are sent depend on the personal momentary context of users (i.e., time and calendar). To examine the usability of such an RL-based smartphone exercise application, a feasibility study with real target users was conducted and presented. We aim to explore both the user behavior (i.e., interactions with the reminders) and the user experience after receiving reminders delivered at the personalized moments. Thus, the following three hypotheses were addressed and examined:

1. Reminders sent at personalized moments can trigger active reactions within a short period of time.
2. Reminders sent at personalized moments can trigger users to start a physical activity.
3. Reminders sent at personalized moments are perceived as useful by users, based on user feedback from questionnaires and interviews.

2. Materials and Methods

To examine the proposed hypotheses, we first designed and integrated an RL model in a smartphone exercise application to deliver reminders based on users' momentary temporal context. Then, we conducted a feasibility study, in which 11 target users participated. At the end, we collected and analyzed data from 7 users who used the smartphone exercise application for 4 weeks and received reminders sent at moments determined by the RL model. The 7 users were also asked to fill out a questionnaire and have an interview to share their user experience.

2.1. Materials

2.1.1. A Smartphone Exercise Application Based on Reinforcement Learning

In our study, a smartphone exercise application (the PAUL app) was developed to send reminders adaptively at moments suggested by a reinforcement learning (RL) model. The development of the PAUL app was conducted based on both theoretical foundations of behavior change [20] and a focus-group study with our target users [21]. The design process and practical implementation of this smartphone exercise application can be found in [22]. In this subsection, we mainly describe the integration of a reinforcement learning model in the PAUL app.

In Figure 1, we present an overview of how the PAUL app iteratively interacts with the user to decide the right moments to deliver reminders, where the RL model is marked in green. In this interaction procedure, the RL model is a decision-maker in the smartphone exercise app. On the one hand, the RL-based decision-maker determines the momentary status of a user based on personal context information collected by the app at each time step, then estimates whether this is the right moment to send a reminder to this user. Furthermore, the RL-based decision-maker records the physical activity behavior of the users after receiving the reminders. In this manner, the RL model learns the preference of this user from his or her historical data, and keeps the state-of-the-art knowledge to make the next decision.

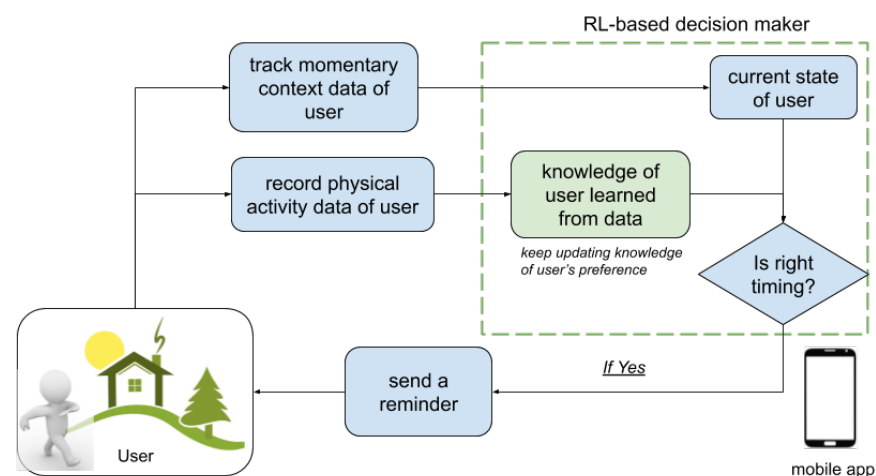


Figure 1. Overview of the reminder delivery procedure in the PAUL smartphone exercise application, regarding integration of a reinforcement learning model.

We decided that the PAUL app should decide on whether to send a reminder or not at every hour from 8:00 to 20:00 in each day. Additionally, since too-frequent interactions with

the user are not desirable in our practical task, we set the maximum number of reminders sent. In general, the PAUL app sends a maximum of 14 reminders each week. The content of the messages were drafted based on previous literature [23] and evaluated in a pilot test. Based on the results of questionnaires, a message library of 141 messages was built. The messages are positively framed, focus on effectiveness and immediate outcomes, and are tailored to the activity type (i.e., running, walking or both).

2.1.2. The Reinforcement Learning Model

We modeled our practical task (i.e., how to learn the optimal strategy for delivering reminders with respect to the momentary context of this user) as a reinforcement learning problem. For each user, we formalized the problem as a Markov decision process (MDP). The MDP framework is a mathematical abstraction of sequential decision-making problems [14]. Figure 2 depicts the interaction between an agent and an environment in our MDP. Here, the agent represents a mobile health system that learns the optimal strategy to interact with a target user, which is our environment. Our agent and environment interact in a sequence of discrete and finite time-steps $\{1, 2, \dots, t\}$, which can be naturally broken into episodes. At each time step, the agent observes the contextual representation of the environment, and on that basis selects an action. Afterwards, the environment passes a numerical reward (inferring its feedback on the given action) back to the agent, and updates itself in a new state. Based on this trial-and-error mechanism, the agent adapts its policy for maximizing an expected long-term reward. We follow the standard definition in [14] and define the key elements of our MDP as follows:

- S : A set of observed *states*, each state $s_t \in S$ is a vector of contextual variables of the environment (i.e., user), including ‘the number of reminders left in this week’, ‘timestep from the last run’, ‘timestep from the last reminder’ and ‘hour of the day’, ‘weekday’, ‘calendar availability’.
- A : A set of *actions* taken by the agent (i.e., app), two actions, ‘sending reminder’ and ‘not sending reminder’, are denoted as $a_t = 1$ and $a_t = 0$ respectively.
- T : A probabilistic transition function, defining the probability of moving from s_{t-1} to s_t given a specific action a_{t-1} .
- R : A *reward* function, representing the probability that a reward r_t was received by the agent. If the target user performs a physical activity at time t , the reward $r_t = 1$; if not, $r_t = 0$.




Figure 2. The interaction flow of our RL model, where the agent representing the smartphone health application and the environment is a user (Please note that the term reward is a specific entity in the MDP model. In our case, it indicates whether a physical activity is performed after the user receives a reminder sent by the RL model).

In our case, we aim to optimize the frequency of physical activities in a week with the PAUL app. Following the setup of sending reminders in the PAUL app, the RL model decides on which action (sending reminder or not) to take at every hour from 8:00 to 20:00 in each day. In total, there are 84 decision time-steps in each week. If a physical activity is performed by the user before the next decision time step (within one hour), we gave a reward of 1.0 to our app (otherwise zero reward is assigned). Since we allow a constrained number of reminders, our optimization goal becomes how to wisely deliver the maximum

number of reminders at the ‘right moments’ in an episode to maximize the user’s physical activity frequency in that period.

2.1.3. The Pre-Learned Reinforcement Learning

In this section, we explain one important property of our RL model. A problem that many RL-based mobile intervention systems suffer from is the cold-start problem [24]. As very few (or even no) experiences with the user are available at the beginning, RL models often require the application to interact many times with the users prior to performing well (learning how to deliver the reminders at the right moments in our case). To remedy this problem, we pre-learned a generalized initial delivery strategy in our RL model from the large-scale empirical running data. An overall framework of our RL smartphone exercise application is shown in Figure 3.




Figure 3. The overall framework for our pre-learned RL exercise application.

Following this procedure, we first clustered users from a large-scale historical running data and identified the target users [25]. The used running data contains around 440,000 runs over 4 years performed by more than 10,000 users. In the data analysis, we characterized the target users as people who struggled with maintaining running behavior (it is consistent with our prioritized participants in the feasibility study). Moreover, as each data point includes the timestamp and weather information at the start moment for every run, we also captured when and under which weather conditions our target users prefer to start runs. Based on those preferences, we then developed an RL algorithm to learn an initial strategy using this empirical data in simulations [26]. In this manner, we aim to obtain a generalized strategy that delivers a limited number of reminders to our target users at reasonable moments (better than a random initialized one in code-start RL models).

2.2. Methods

2.2.1. Study Design

The feasibility study consisted of three steps, namely enrollment, intervention and feedback. First, in the enrollment step, we recruited the participants, visited all eligible participants and asked them to fill out a questionnaire to assess their socio-cognitive characteristics. Second, our participants used the PAUL smartphone application in their own smartphone from 18 November to 15 December 2019. They were asked to turn on the reminder function of their smartphones (i.e., to receive pop-up notifications) and to give access to their digital calendar (i.e., to provide calendar availability data, but they did not have to). We separated this 4-week study period into 2 phases: baseline (3 weeks) and intervention (1 week). In the first 3 weeks, participants had access to the PAUL app, but no JITAI reminders were sent by the app. These initial three weeks without the JITAI reminders allowed us to ensure the usability of basic functions in the app and to collect the baseline behaviors of participants. Then, in the final intervention week, the PAUL app delivered a maximum of 14 reminders to each participant, asking them to perform a physical activity (either a walk or a run). The moments of delivery were decided by the RL model based on the user’s personal momentary information (i.e., time preference and

calendar availability). During the whole period, all participants were encouraged (but not obliged) to use the PAUL app as much as possible, and several data on participants' behaviors (e.g., interactions with the app) were monitored and collected automatically by the PAUL app. Third, at the end of the study, all participants answered a questionnaire on the usability of the reminder function in the PAUL app. In this feedback step, we also conducted a one-to-one interview (20–40 min) with each participant, to understand their experience and feelings on using the PAUL app in the manner we designed.

2.2.2. Participant Recruitment

After being granted approval from the local ethics review board (No. ERB Review Geo S-19253), we started our study by recruiting participants. To maintain a healthy lifestyle, people are advised to engage in a sufficient amount of physical activity on a regular basis. Therefore, our prioritized participants are people who are struggling with maintaining a healthy activity level and would like to participate in more physical activities. During the recruitment, we also accepted participants who perform physical activity regularly but still need to move more to reach the suggested amount.

To control the influence of living location on user behavior, we selected three parks in city centers of the Netherlands (i.e., Transwijk at Utrecht, Oosterpark at Amsterdam and Slotterplas at Amsterdam) and only recruited participants who live close (less than 5 km and 20 min bike ride) to either park. We built up a website to introduce the relevant information of our study for interested individuals and to include an online eligibility questionnaire, which is used for assessing their stage of healthy behavior change (as defined by the Transtheoretical Model (TTM) [27]). Our stakeholders in the project also aided with the recruitment by contacting neighborhood organizations (e.g., promotion materials such as flyers were placed in public areas around the parks). We also recruited participants via Facebook advertisements. Based on the participant applications, we enrolled 11 participants who were aged between 18 and 55 years and met the following criteria: not meeting the physical activity guidelines of 150 min per week, with no medical condition preventing them from performing PA (defined by PAR-Q [28]), owning an Android smartphone, not currently participating in another PA or health-related intervention and have proficient knowledge of the Dutch language (as the app is in Dutch). In total, 7 of 11 participants completed the whole study and received reminders sent at personalized moments by the RL model, whose demographic information is presented in Table 1.

Table 1. The demographical information of 7 participants who completed the feasibility study.

Type	Category	Value
Gender		71.4% Female, 28.6% Male
Age		At average 34.4 (range from 21 to 51)
BMI		At average 24.73 (range from 19.72–36.99)
Residence	Near Utrecht Transwijk	3
	Near Amsterdam Oosterpark	3
	Near Amsterdam Slotterplas	1
Education level	Secondary school (VWO)	1
	Vocational education	1
	Associate degree	1
	University degree	4
Employment status	Part-time employment (<34 h/w)	3
	Full-time employment (≥34 h/w)	1
	Studying	3
Stage of change (moderate activity)	Maintenance stage	4
	Action stage	1
	Contemplation stage	2

2.2.3. Outcome Measurements, Questionnaire and Interview

During the study, our PAUL smartphone app automatically tracked the physical activities performed by the participants and recorded their reactions after receiving reminders. On the one hand, the data collection of physical activity starts when the user clicks the ‘start exercise’ button of the app and then continuously tracks all the data involved in a run in the background, until it is terminated by the user in a comparable manner. For each physical activity, a dataset is collected summarizing total distance and run time, and marking the timestamp and GPS location at the start point of this physical activity. On the other hand, the data collection of users’ active reactions to reminders starts when they give positive feedback. The positive feedback can be given by either tapping on the pop-up notification corresponding to the reminder or opening the application in a short period of time. Otherwise, we consider users ignore the reminders because they were not sent at the right moments. Based on the collected data, we address the following two outcome measurements.

- **The number of reminders actively reacted by a user:** a user was considered to ‘actively react to a reminder’ if he or she tapped on the pop-up notification after this reminder was delivered.
- **The number of physical activities triggered by a reminder:** a physical activity was defined as ‘triggered by a reminder’ if the user performed a physical activity after he or she actively reacted to a delivered reminder.

Using those measurements, we aim to quantify user behavior after receiving our designed reminders and examine the first two hypotheses in Section 1.

In addition, to explore the experience of users on receiving the reminders and to answer the third hypothesis, we designed several questions for a questionnaire and an individual interview based on previous studies [13,29,30]. These questions also concentrate on the receptivity of reminders, which is a concept that anticipates a subjective overall reaction of the user to a notification [30]. We particularly measure two underlying factors of receptivity, which are the willingness to be interrupted (known as reachability) [29] and the likelihood of influencing the recipient’s future actions (known as actionability) [30]. All interviews were audio-recorded and transcribed. In Table 2, we present the questions asked in the questionnaire and the interview respectively (the original ones are all in Dutch). Among them, all questionnaire questions are multiple choices with a 4- or 5-point Likert scale (see the scale details in Section 3.2).

Table 2. Questions in the questionnaire and the interview about user experience on received reminders.

Questions in the Questionnaire
1. How often did you think the reminder was sent at a good time?
2. How many times had the reminders moved you with the app?
3. How often did you want to move with the app when you received the reminder?
4. How often did you immediately notice the reminder (when it had just been sent)?
5. How often did you feel interrupted by the reminders?
6. How often did you feel annoyed by the reminders?
7. Have you ever turned off the reminder function (by turning off the push notifications)?
8. What do you think of the number of reminders you received?
9. How much time would you like to have between the reminder and your exercise session?
Topic list of the interview
1. What do you think of the reminders sent by the app?
2. Did you find it useful? And why?
3. What did you think of the number of reminders? What would be your ideal number of reminders?
4. Were you satisfied with the timing of the reminders? Did the reminders bother you? Why? Do you have any suggestions for changing the timing?
5. Did the reminders motivate you to move more? How? Why or why not?
6. Can you think of something to improve the feature that we haven’t discussed yet?

2.2.4. Data Analysis

In this study, we mainly used descriptive analysis and analyzed the data collected from three resources (including the PAUL application, a questionnaire, and interviews). We not only statistically described user objective behaviors from the PAUL application data, but also conducted the descriptive analysis to understand user experience from the questionnaire and interview data.

Regarding the PAUL application data, two datasets were collected and analyzed. On the one hand, a log dataset collected the app usage events of every participant and particularly marked the timestamp of those events. Table 3 gives examples of our collected data in the log dataset. Following two defined outcome measurements in Section 2.2.3 (i.e., reminders actively reacted to by a user and physical activity triggered by a reminder), we tracked the objective behaviors of users after receiving reminders from the PAUL application. On the other hand, a session dataset contained detailed information about all physical activity sessions (i.e., user ID, starting timestamp, duration, length, etc.).

Table 3. Example data records indicate the behaviors of a user after receiving a reminder.

Event ID	User ID	Timestamp	Activity
3535	NO. 2	11 December 2019 15:00:00	reminder sent
3536	NO. 2	11 December 2019 15:04:00	notification clicked
3537	NO. 2	11 December 2019 15:04:01	app opened
3538	NO. 2	11 December 2019 15:07:00	walking or running session started

During this process, to make sure the data analysis is valid, we confirmed the consistency of collected data. For instance, a walking or running session data entity was included in the analysis only if its starting timestamp was marked as the same in both datasets from the PAUL app. Additionally, to determine the reliability of the user experience data, we conducted a cross-check between answers in the questionnaire and interviews of same participants.

3. Results

In this section, we present the results of the usability of our RL-based application, including user reaction times and physical activity behaviors, as well as their perceived feedback.

3.1. Results Related to User's Objective Behaviors

In this section, we present the quantitative results based on the quantified outcome measurements defined in Section 2.2.3, which are mainly related to users' objective behaviors.

First, we summarized the physical activities performed by our participants, where the records with less than 60 s were excluded (i.e., the duration of run or walk is less than one minute). We realized that 6 out of 20 physical activities performed by all 7 users were in the intervention week (i.e., accounting for 30%), which is higher than the average in all four weeks. Meanwhile, in total, 79 reminders were sent to our participants in the intervention week. On average, each user received 11.3 reminders per week.

Meanwhile, we present the results of two kinds of reactions performed by all users in Figure 4. We observed that participants tend to behave in a diverse manner after receiving the designed reminders. For instance, user No. 3 actively clicked on most of the pop-up reminders sent to him or her, while none of them transferred into an actual physical activity performance. Conversely, although user Nos. 2, 4 and 7 only reacted to a few reminders (by clicking them), they did seem to be triggered by those reminders and performed physical activities afterwards. To demonstrate the detailed behavior of participants, we present the behavior of user No. 2 during the intervention week in Figure 5 (who was the most active participant in our study).




Figure 4. The reaction behaviors of all 7 users after receiving reminders from the PAUL app.

	Mon	Tue	Wed	Thu	Fri	Sat	Sun
8:00							
9:00				🔔	🔔		
10:00		🔔	🔔	🔔	🔔		
11:00							🔔✓
12:00							
13:00			🔔✓				
14:00		😊					
15:00			🔔✓😊				🔔✓😊
16:00					🔔	🔔	
17:00							
18:00					🔔		
19:00							

Figure 5. The reminder delivery and corresponding behaviors of user No. 2 in the intervention week, where a bell represents one reminder, a tick represents one short-term reaction (clicking the reminder pop-up), and a smile represents one physical activity.

In addition, by summarizing the data of our participants, we noticed several common behaviors. First, we observed that 4 out of 6 physical activities in this week were performed within 5 h of the delivery of a reminder (accounting for about 66.7%). Those 4 physical activities were contributed by 3 different participants (i.e., user Nos. 2, 4 and 7). Interestingly, user No. 7 only ran one time during the whole feasibility study, which was triggered by our reminder. Second, we found that 18 out of 79 reminders received a reaction from our participants by clicking on the pop-up notification (accounting for 22.8%). Importantly, every participant reacted to at least one reminder. The distribution of the interval time between the reaction (timestamp to click a reminder) and delivery of a reminder is calculated and given in Figure 6. From the figure, we can observe that most reminders (83.3%) were actively reacted to by our participants within 50 min.

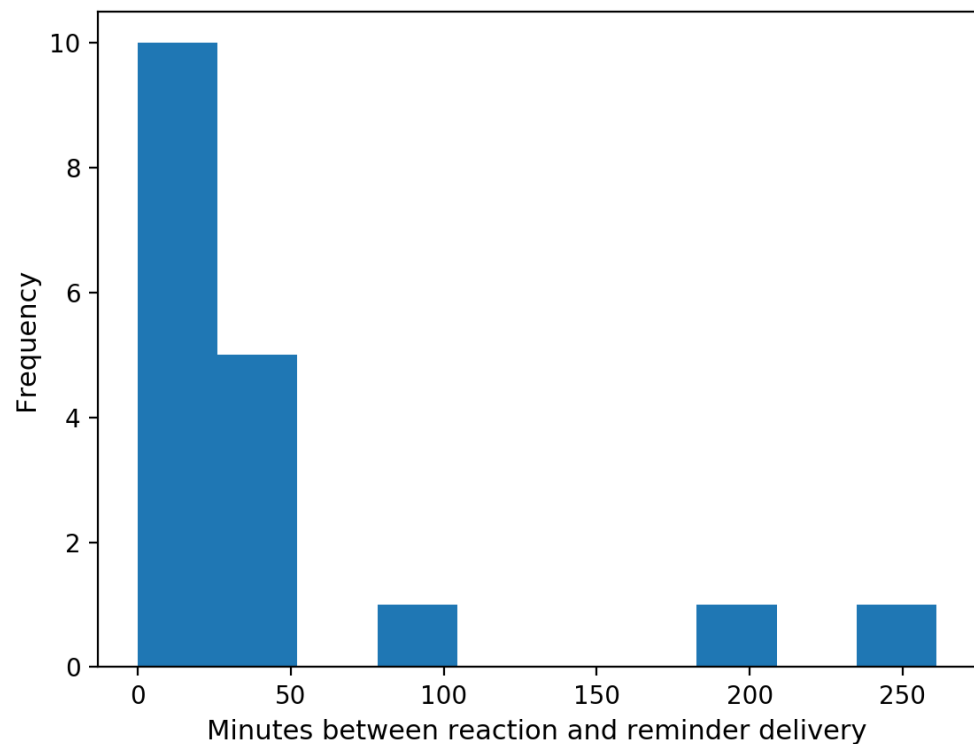


Figure 6. Distribution of the interval time between reminder delivery and reaction (when participants clicked a reminder).

3.2. Results of Questionnaire and Interview

In this section, we present the qualitative results based on a questionnaire and an interview. We separate the detailed results of the questionnaire in Figure 7 (Q1–Q6 related to the timing of the reminders) and Table 4 (Q7–Q9 related to other subjective experience of the reminders), while we jointly discuss the findings from the questionnaire and the interview.

In summary, all participants indicated that it was very important to receive reminders to initiate a running or walking session. Some explained that reminders ‘confront’ them with not doing the exercise and reaching their goal (user No. 5) or that they would otherwise forget to go for a run or walk (user No. 7) or to use the application (user No. 3). One of the participants also claimed that although the reminders of applications sometimes annoyed her, she did think they are useful, because you would not download such an application if you do not need to be pushed:

“...no signal at all is not a kick in the ass, and apparently you need it, that’s why you need an app, yeah, and you install it” (user No. 2)

Furthermore, our participants also highlighted that the right ‘timing’ of reminders is significant to make sure they are effective. If the messages are continuously sent at times that people cannot engage in physical activity, they can end up feeling annoyed, or even feel discouraged and disappointed in themselves for not being able to do more activity, while they do want to do it:

“And every time you think, “Now I can’t, now I can’t.” [sighs] And that feels like giving up. [...] That’s just demotivating for, um, using the app very much. And not necessarily for moving. That you think after a while: “Oh I remove the app, because this is not nice anymore”. Like that. I didn’t have that yet, but it can happen.” (user No. 7)

According to Q4 in Figure 7, more than half of the participants ($n = 4$) noticed most reminders immediately after their delivery (including 2 participants almost always noticed them). This result is consistent with the quantified results of number of reminders actively reacted by users shown in Figure 6, which indicates that our reminders were sent at the right timing of ‘capturing the immediate attention of a user’.




Figure 7. The results of the top 6 questions in the questionnaire, showing perceptions of the timing of reminders. Each unit on the X coordinate represents one participant.

Table 4. The results of other questions in the questionnaire about reminder-related properties.

Question	Choices	Value
Q7. Have you ever turned off the reminder function, by turning off the push notification?	Yes, (almost) the entire intervention period	0
	Yes, for multiple days	1
	Yes, but for a day or multiple hours	0
	No	6
Q8. What do you think of the number of reminders you received?	way too few	2
	too few	1
	good	2
	too much	2
	way too much	0
Q9. How much time would you like to have between the reminder and your exercise session?	less than 10 min	0
	less than 1 h	2
	about 1 h	4
	between 1 and 2 h	0
	longer than 2 h	1

In addition, we found that three participants did not perceive the reminders to be sent at the right times (see Q1 in Figure 7). Two of them also stated that there were too few messages (one participant even claimed that he or she never received any reminders), which might explain their negative feedback about the delivery timing. In other words, there is a possibility that the reminders were not sent at times they needed them (e.g., in the first 3 weeks or as frequent as they wanted), instead of at bad times. We also noticed that our reminders did not always motivate people to perform a physical activity. Only about 57% of participants (n = 4) perceived the reminders can motivate them to engage in physical activity (see Q2 in Figure 7), which is similar to the number of participants who actually performed a physical activity during the intervention week (n = 3) (see Figure 4). Interestingly, some participants discussed the reasons that they thought the reminders were not sent in good times in the interview. One important factor is the weather. Since the Dutch weather tended not to be attractive during our feasibility study (e.g., windy, rainy and short daytime), a few participants claimed that they were not very motivated to perform in such weather conditions.

Most participants did not feel annoyed (n = 4) or interrupted (n = 4) by the reminders (see Q5 and Q6 in Figure 7), while they perceived the reminder function (e.g., only one participant turned off the reminders for a couple of days in Table 4). Thus, although the reminders were not always sent at optimal times, they also did not frustrate people to a degree that would result in stopping using the application or perform physical activity.

During the interviews, different topics arose regarding the frequency of the reminders. In general, it seems the preferred number of reminders is related to the goal of the participants. Some participants want to plan their activities and receive a certain number of reminders based on their own settings. However, other people, who do not like to plan

their activities, prefer to have reminders daily, so they can decide at that point of time whether they feel like running or walking. See the following answers from the participants:

“Well, if I set a goal of three times a week, I want a reminder three times a week. And that I can choose on which days, so I can schedule it in advance.” (user No.5)

“Um, I think one every day, because otherwise a day will soon go by. Say with one of those apps, I’ll miss it quickly. So if there is such a reminder, I thought “oh yeah, I should do that too”. And if you don’t, after all, chances are I won’t do anything about it that day. And if I got that reminder, I’m more likely to think ‘oh yes, I had this as my goal, I have to do this for a while’.” (user No.3)

4. Discussion and Conclusions

In this paper, we follow just-in-time adaptive intervention (JITAI) and propose to send reminders at adaptive times based on users’ momentary temporal and calendar information. We first integrated the reinforcement learning approach in a mobile exercise application for determining the right moments to deliver a reminder. Then, we conducted a 4-week feasibility study and evaluated the usability of such a reminder delivery method. This study demonstrates that sending reminders at personalized moments through a RL-based mobile application is feasible. These JITAI reminders were sent at times that the participants appear to be receptive, but not always at times that the participants had the opportunity to exercise.

First, the results shown in Section 3 provided positive evidence for the hypotheses. We found that reminders sent in our personalized moments were able to attract user reactions within a short period of time. For instance, all participants actively reacted to the reminders and over 83.3% reactions were done within 50 min of receiving the reminders. This is an important property to increase the reachability of reminders [31] (introduced in Section 2.2.3), as it can respect the user’s willingness of being interrupted. For instance, if a participant is busy with something, it is hard to interrupt him and let him receive the content of this reminder. Meanwhile, reminders sent in our personalized moments could trigger users to start a physical activity. For instance, an inactive participant only performed one physical activity during the entire feasibility study, which was triggered by our reminder. In the interviews, our participants also highlighted that the right ‘timing’ of reminders are significant to make sure they are effective. Those findings are in line with earlier research that indicated the effectiveness of delivery ‘timing’ in mobile interventions [7,8].

Several interesting scientific contributions have also been made in our study. First, it is among the first studies to examine the relationship between the reminder time, the reaction time and the action time (actual performance time) for physical activity interventions. For instance, according to the questionnaire, most participants would like to have about 1 h between the reminder and the actual exercise. In addition, shown in the behavioral results, the reaction time and action time were about 1 h and 5 h respectively. As the reaction and action behaviors could reflect the reachability and actionability of interventions respectively [29,30], we think those two interval times (action time and reaction time) shall be further explored in determining the effectiveness of mobile health interventions. Second, as far as we know, our RL method incorporated a constraint on interaction frequency (i.e., the maximum number of reminders per week) for the first time in RL-based mobile intervention applications. According to the questionnaires and interviews, only 2 out of 7 users expressed that there were too many reminders. Additionally, most of participants rarely felt annoyed or interrupted by the reminders. As the frequency of interactions could add burden for the user engagement [19], we therefore recommend that future applications also restrict the frequency of reminders when developing RL methods for mobile health interventions. Furthermore, we generated results from both quantified outcome measurements and qualitative studies (i.e., questionnaire and interview). By comparing those results, we noticed an interesting phenomenon. A participant claimed that he or she did not receive any reminder in the interview (in Section 3.2); however, based on our data collection, all participants have clicked on the notification of a reminder at least once

(see Figure 4). We think this phenomenon was caused by the deviation between physical behavior and awareness of the participant. As discussed in a recent review, the automatic detection from sensing data could reduce the burden of self-reporting [32]. This finding indicated the importance of combining results from the non-invasive methods with the qualitative methods in feasibility studies.

Besides the interesting findings, there are still several limitations to our research. First, a general methodological challenge concerns the fact that this is only a small-scale feasibility study, where the number of participants was limited, and the duration of intervention was relatively short. These setups could influence the accuracy of our results. Additionally, due to the short duration, our RL model only used a pre-learned delivery for sending the reminders and was not able to adjust the strategy for each individual participant. Those reasons drive us to conduct a large-scale study with more participants and a longer intervention period. As examining the feasibility prior to large-scale effectiveness testing is an important preparation step in the development phase of a digital intervention [33], we think the study protocol and lessons learned in this feasibility study will be useful in the future study. Additionally, to examine the effectiveness of our reminder delivery function, we would like to cover a comparison between the adaptive timing and the non-adaptive timing in the large-scale study. Second, although we did not set up any gender bias during the participant recruitment, we are aware that most of our participants were female (accounting for 71.4%). Furthermore, due to the technical restriction, only people who have an Android smartphone could participate in this study. For those two reasons, it may be the case that our results are not generalized enough. We therefore would like to overcome this problem in the future by better balancing the gender of participants and expanding the possible users.

The lessons learned from this study also motivate us to improve our intervention design in future research. For instance, as shown from the user experience results, we still need to improve the delivery timing of our reminders for influencing the user's future actions, which corresponds to the actionability of reminders [30]. As mentioned by some participants, factors such as the weather can be very effective. Previous JITAI studies also indicated that several factors (e.g., weather and personal mood) can play an important role in determining well personalized timing of reminders [5]. We therefore would like to take more momentary context information into account to improve the strategy of our reinforcement learning approach. Second, similar to the preferred timing, the preferred frequency of reminders is dramatically different among participants based on the questionnaire and interview results. For instance, some participants prefer 2 or 3 messages a week, while others would like to have two a day. We think this phenomenon will be taken into account in the future study. One possible solution is to give a different initialized strategy to the different user clusters based on their frequency preference [34].

Author Contributions: Conceptualization, K.S. and S.W.; methodology, S.W. and H.v.H.; smartphone app development, R.D.D.d.B. and S.W.; feasibility study, K.S., S.W., M.S. and N.N.; data analysis, S.W.; interview, K.S.; writing—original draft preparation, S.W., K.S. and B.K.; writing—review and editing, all authors; supervision, M.S., D.E. and B.K.; funding acquisition, M.D. and B.K. All authors have read and agreed to the published version of the manuscript.

Funding: This work is funded by Playful Data-driven Active Urban Living project under NWO and SIA grant 629.004.013.

Institutional Review Board Statement: The study was conducted according to the guidelines of the Declaration of Helsinki, and approved by the Science-Geo Ethics Review Board at Utrecht University (Geo S-19253, 26 September 2019).

Informed Consent Statement: Informed consent was obtained from all subjects involved in the study.

Data Availability Statement: The datasets for this article are not publicly available because the collected data (e.g., application usage data, interviews) are privacy sensitive. No consent has been given to publicly share this data, but the data can be made available on request for verification purposes. Request to access the datasets should be directed to the corresponding author (s.wang2@uu.nl).

Acknowledgments: The authors especially thank the cooperator MYLAPS for providing the smartphone application dataset. We also appreciate the precious comments from our colleagues Victor Dourado, Paula Castro and Chao Zhang.

Conflicts of Interest: The authors declare no conflict of interest.

References

- Warburton, D.E.; Nicol, C.W.; Bredin, S.S. Health benefits of physical activity: The evidence. *CMAJ* **2006**, *174*, 801–809. [[CrossRef](#)] [[PubMed](#)]
- Weggemans, R.M.; Backx, F.J.; Borghouts, L.; Chinapaw, M.; Hopman, M.T.; Koster, A.; Kremers, S.; van Loon, L.J.; May, A.; Mosterd, A.; et al. The 2017 Dutch physical activity guidelines. *Int. J. Behav. Nutr. Phys. Act.* **2018**, *15*, 1–12. [[CrossRef](#)]
- Reis, R.S.; Salvo, D.; Ogilvie, D.; Lambert, E.V.; Goenka, S.; Brownson, R.C.; Lancet Physical Activity Series 2 Executive Committee. Scaling up physical activity interventions worldwide: Stepping up to larger and smarter approaches to get people moving. *Lancet* **2016**, *388*, 1337–1348. [[CrossRef](#)]
- Griffiths, F.; Lindenmeyer, A.; Powell, J.; Lowe, P.; Thorogood, M. Why are health care interventions delivered over the internet? A systematic review of the published literature. *J. Med. Internet Res.* **2006**, *8*, e10. [[CrossRef](#)]
- Rabbi, M.; Aung, M.H.; Zhang, M.; Choudhury, T. MyBehavior: Automatic personalized health feedback from user behaviors and preferences using smartphones. In Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing, ACM, Osaka, Japan, 7–11 September 2015; pp. 707–718.
- Hardeman, W.; Houghton, J.; Lane, K.; Jones, A.; Naughton, F. A systematic review of just-in-time adaptive interventions (JITAs) to promote physical activity. *Int. J. Behav. Nutr. Phys. Act.* **2019**, *16*, 31. [[CrossRef](#)]
- Nahum-Shani, I.; Smith, S.N.; Spring, B.J.; Collins, L.M.; Witkiewitz, K.; Tewari, A.; Murphy, S.A. Just-in-time adaptive interventions (JITAs) in mobile health: Key components and design principles for ongoing health behavior support. *Ann. Behav. Med.* **2017**, *52*, 446–462. [[CrossRef](#)]
- Fogg, B.J. Persuasive technology: Using computers to change what we think and do. *Ubiquity* **2002**, *2002*, 2. [[CrossRef](#)]
- Rahman, T.; Czerwinski, M.; Gilad-Bachrach, R.; Johns, P. Predicting “about-to-eat” moments for just-in-time eating intervention. In Proceedings of the 6th International Conference on Digital Health Conference, Montreal, QC, Canada, 11–13 April 2016; pp. 141–150.
- Goldstein, S.P.; Evans, B.C.; Flack, D.; Juarascio, A.; Manasse, S.; Zhang, F.; Forman, E.M. 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*, 673–682. [[CrossRef](#)] [[PubMed](#)]
- Sarker, H.; Sharmin, M.; Ali, A.A.; Rahman, M.M.; Bari, R.; Hossain, S.M.; Kumar, S. Assessing the availability of users to engage in just-in-time intervention in the natural environment. In Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing, Seattle, WA, USA, 13–17 September 2014; pp. 909–920.
- Thomas, J.G.; Bond, D.S. Behavioral response to a just-in-time adaptive intervention (JITAI) to reduce sedentary behavior in obese adults: Implications for JITAI optimization. *Health Psychol.* **2015**, *34*, 1261. [[CrossRef](#)]
- Ding, X.; Xu, J.; Wang, H.; Chen, G.; Thind, H.; Zhang, Y. WalkMore: Promoting walking with just-in-time context-aware prompts. In Proceedings of the 2016 IEEE Wireless Health (WH), Bethesda, MD, USA, 25–27 October 2016; pp. 1–8.
- Sutton, R.S.; Barto, A.G. *Reinforcement Learning: An Introduction*; MIT Press: Cambridge, MA, USA, 2018.
- Yom-Tov, E.; Feraru, G.; Kozdoba, M.; Mannor, S.; Tennenholtz, M.; Hochberg, I. Encouraging physical activity in patients with diabetes: Intervention using a reinforcement learning system. *J. Med. Internet Res.* **2017**, *19*, e338. [[CrossRef](#)]
- Zhou, M.; Mintz, Y.; Fukuoka, Y.; Goldberg, K.; Flowers, E.; Kaminsky, P.; Castillejo, A.; Aswani, A. Personalizing Mobile Fitness Apps using Reinforcement Learning. In Proceedings of the IUI Workshops, Tokyo, Japan, 7–11 March 2018.
- Forman, E.M.; Kerrigan, S.G.; Butryn, M.L.; Juarascio, A.S.; Manasse, S.M.; Ontañón, S.; Dallal, D.H.; Crochiere, R.J.; Moskow, D. Can the artificial intelligence technique of reinforcement learning use continuously-monitored digital data to optimize treatment for weight loss? *J. Behav. Med.* **2019**, *42*, 276–290. [[CrossRef](#)]
- Liao, P.; Greenewald, K.; Klasnja, P.; Murphy, S. Personalized HeartSteps: A Reinforcement Learning Algorithm for Optimizing Physical Activity. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* **2020**, *4*, 1–22. [[CrossRef](#)]
- Pellegrini, C.A.; Pfammatter, A.F.; Conroy, D.E.; Spring, B. Smartphone applications to support weight loss: Current perspectives. *Adv. Health Care Technol.* **2015**, *1*, 13. [[CrossRef](#)]
- Michie, S.; Van Stralen, M.M.; West, R. The behaviour change wheel: A new method for characterising and designing behaviour change interventions. *Implement. Sci.* **2011**, *6*, 42. [[CrossRef](#)]
- Nibbeling, N.; Simons, M.; Sporrel, K.; Deutekom, M. Perceptions of inactive adults regarding a theory-based physical activity app: A focus group study. *Front. Public Health* **2021**. [[CrossRef](#)]

22. Sporrel, K.; De Boer, R.D.D.; Wang, S.; Nibbeling, N.; Simons, M.; Deutekom, M.; Ettema, D.; Castro, P.C.; Dourado, V.Z.; Kröse, B. The Design and Development of a Personalized Leisure Time Physical Activity Application Based on Behavior Change Theories, End-User Perceptions, and Principles From Empirical Data Mining. *Front. Public Health* **2021**, *8*, 711. [[CrossRef](#)] [[PubMed](#)]
23. Latimer, A.E.; Brawley, L.R.; Bassett, R.L. A systematic review of three approaches for constructing physical activity messages: What messages work and what improvements are needed? *Int. J. Behav. Nutr. Phys. Act.* **2010**, *7*, 36. [[CrossRef](#)]
24. Tabatabaei, S.A.; Hoogendoorn, M.; van Halteren, A. Narrowing reinforcement learning: Overcoming the cold start problem for personalized health interventions. In Proceedings of the International Conference on Principles and Practice of Multi-Agent Systems, Tokyo, Japan, 29 October–2 November 2018; pp. 312–327.
25. Wang, S.; Scheider, S.; Sporrel, K.; Deutekom, M.; Timmer, J.; Kröse, B. What Are Good Situations for Running? A Machine Learning Study Using Mobile and Geographical Data. *Front. Public Health* **2021**, *8*, 985. [[CrossRef](#)] [[PubMed](#)]
26. Wang, S.; Zhang, C.; Kröse, B.; van Hoof, H. Optimizing Adaptive Notifications in Mobile HealthInterventions Systems: Reinforcement Learning from a Data-driven Behavioral Simulator. **2021**, under review.
27. Prochaska, J.O.; Velicer, W.F. The transtheoretical model of health behavior change. *Am. J. Health Promot.* **1997**, *12*, 38–48. [[CrossRef](#)]
28. Thomas, S.; Reading, J.; Shephard, R.J. Revision of the physical activity readiness questionnaire (PAR-Q). *Can. J. Sport Sci.* **1992**, *17*, 338–345. [[PubMed](#)]
29. Begole, J.B.; Matsakis, N.E.; Tang, J.C. Lilsys: Sensing unavailability. In Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work, Chicago, IL, USA, 6–10 November 2004; pp. 511–514.
30. Fischer, J.E.; Yee, N.; Bellotti, V.; Good, N.; Benford, S.; Greenhalgh, C. Effects of content and time of delivery on receptivity to mobile interruptions. In Proceedings of the 12th International Conference on Human Computer Interaction with Mobile Devices and Services, Lisbon, Portugal, 7–10 September 2010; pp. 103–112.
31. Pejovic, V.; Musolesi, M. InterruptMe: Designing intelligent prompting mechanisms for pervasive applications. In Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing, Seattle, WA, USA, 13–17 September 2014; pp. 897–908.
32. Thomas Craig, K.J.; Morgan, L.C.; Chen, C.H.; Michie, S.; Fusco, N.; Snowdon, J.L.; Scheufele, E.; Gagliardi, T.; Sill, S. Systematic review of context-aware digital behavior change interventions to improve health. *Transl. Behav. Med.* **2020**, 1037–1048. [[CrossRef](#)]
33. Mummah, S.A.; Robinson, T.N.; King, A.C.; Gardner, C.D.; Sutton, S. IDEAS (Integrate, Design, Assess, and Share): A framework and toolkit of strategies for the development of more effective digital interventions to change health behavior. *J. Med. Internet Res.* **2016**, *18*, e317. [[CrossRef](#)] [[PubMed](#)]
34. el Hassouni, A.; Hoogendoorn, M.; van Otterlo, M.; Barbaro, E. Personalization of health interventions using cluster-based reinforcement learning. In Proceedings of the International Conference on Principles and Practice of Multi-Agent Systems, Tokyo, Japan, 29 October–2 November 2018; pp. 467–475.