



OPEN ACCESS

EDITED BY

Chao Zhang,
Eindhoven University of Technology,
Netherlands

REVIEWED BY

Varun Mishra,
Northeastern University, United States
Sian Joel-Edgar,
Northeastern University, United Kingdom

*CORRESPONDENCE

Radoslava Švihrová

✉ radoslava.svihrova@gmail.com;

✉ radoslava.svihrova@unibe.ch

RECEIVED 21 May 2024

ACCEPTED 21 May 2025

PUBLISHED 05 June 2025

CITATION

Švihrová R, Dei Rossi A, Marzorati D, Tzovara A and Faraci FD (2025) Designing digital health interventions with causal inference and multi-armed bandits: a review.
Front. Digit. Health 7:1435917.
doi: 10.3389/fdgth.2025.1435917

COPYRIGHT

© 2025 Švihrová, Dei Rossi, Marzorati, Tzovara and Faraci. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](#). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Designing digital health interventions with causal inference and multi-armed bandits: a review

Radoslava Švihrová^{1,2*}, Alvise Dei Rossi^{2,3}, Davide Marzorati²,
Athina Tzovara^{1,4} and Francesca Dalia Faraci²

¹Institute of Computer Science, Faculty of Science, University of Bern, Bern, Switzerland, ²Institute of Digital Technologies for Personalized Healthcare, Department of Innovative Technologies, University of Applied Sciences and Arts of Southern Switzerland, Lugano, Switzerland, ³Faculty of Informatics, Università Della Svizzera Italiana, Lugano, Switzerland, ⁴Center for Experimental Neurology, Department of Neurology, Inselspital, Bern University Hospital and University of Bern, Bern, Switzerland

Recent statistics from the World Health Organization show that non-communicable diseases account for 74% of global fatalities, with lifestyle playing a pivotal role in their development. Promoting healthier behaviors and targeting modifiable risk factors can significantly improve both life expectancy and quality of life. The widespread adoption of smartphones and wearable devices enables continuous, in-the-wild monitoring of daily habits, opening new opportunities for personalized, data-driven health interventions. This paper provides an overview of the advancements, challenges, and future directions in translating principles of lifestyle medicine and behavior change into AI-powered mobile health (mHealth) applications, with a focus on Just-In-Time Adaptive Interventions. Considerations for the design of adaptive interventions that leverage wearable and contextual data to dynamically personalize behavioral change strategies in real time are discussed. Bayesian multi-armed bandits from reinforcement learning are exploited as a framework for tailoring interventions, with causal inference methods used to incorporate structural assumptions about the user's behavior. Furthermore, strategies for evaluation at both individual and population levels are presented, with causal inference tools to further guide unbiased estimates. A running example of a simple real-world scenario aimed at increasing physical activity through digital interventions is used throughout the paper. With input from domain experts, the proposed approach is generalizable to a wide range of behavior change use cases.

KEYWORDS

behavioral change, causal inference, digital health, interventions, micro-randomized trials, multi-armed bandit, personalized medicine, statistical analysis

1 Introduction

Non-communicable diseases (NCDs) are a challenge to global public health (1). WHO statistics show that NCDs account for a substantial majority (74%) of worldwide mortality (1). Among the leading contributors to this burden are cancer, diabetes, and cardiovascular diseases, each linked to lifestyle factors such as diet, physical activity, and tobacco use (2). In recent years, there has been a growing recognition of the potential

for lifestyle interventions to mitigate the prevalence of NCDs and improve overall health outcomes (3).

With the global wide adoption of mobile and wearable devices and associated applications, lifestyle management through health apps is becoming more and more accessible to a broader population (4). Emerging technologies offer novel opportunities for real-time monitoring, personalized feedback, and intervention delivery, thereby empowering individuals to take charge of their health in modern ways and in clinical practice (5, 6). In parallel, the fast raising field of artificial intelligence (AI) holds extensive promise for enhancing the effectiveness and scalability of behavior change interventions (7–9). By harnessing the power of machine learning algorithms, AI can analyze vast amounts of data to derive insights from individual behaviors, preferences, and contexts, thus enabling the delivery of tailored interventions that resonate with users on a personal level (10).

The present work reviews the current perspectives and challenges in both designing and evaluating the mHealth behavior change solutions. After introducing the concepts of lifestyle medicine and behavioral change, we address how to integrate them in behavioral change support systems that leverage personalization techniques, including just-in-time adaptive interventions. Further, usage of bayesian multi-armed bandits to enhance user engagement and adherence to lifestyle change programs is discussed. In the second part of the paper, current evaluation techniques for interventional data are presented, with focus on sequential designs. Finally, recent challenges and potential developments in the field are described. Throughout this paper, we will consider as a use case scenario the design of a digital intervention system aimed at improving the level of *physical activity*, which represents one of the fundamental lifestyle pillars. The discussed framework is however generalizable, after input from experts in the field, to any of the other lifestyle pillars or their combination.

1.1 Lifestyle medicine

Incorporating principles of lifestyle medicine into a mobile app for interventions is essential for promoting holistic health and well-being. Lifestyle medicine is a branch of evidence-based medicine focusing on comprehensive behavioral changes through interventions aiming to prevent, manage, and sometimes even treat various health conditions (11). The role of lifestyle factors in influencing health outcomes is emphasized. According to (12) there are 6 pillars of lifestyle medicine recognizing impact on overall health. These are namely: mental well-being and stress-management; healthy relationships; healthy eating; physical activity; sleep; minimizing harmful substances.

The aim of lifestyle medicine is to address the root causes of health issues by promoting healthy behaviors and habits (11). Lifestyle medicine is often integrated into conventional medical practices and may involve collaboration among healthcare professionals, including physicians, dietitians, exercise physiologists, psychologists, and other specialists. It's a holistic

approach that considers the interconnectedness of various lifestyle factors and their impact on health.

1.1.1 Behavioral change

Comprehending the basics of psychology related to behavior change is imperative when creating an app for an intervention framework (13). According to the transtheoretical model of health behavior change (14) there are 6 stages of behavioral change: precontemplation, contemplation, preparation, action, maintenance and termination. In precontemplation stage a need of change is formulated, followed by contemplation stage with motivation and definition of distal (long-term) goals and plans. During the preparation stage, the individual sets realistic proximal (short-term) goals and gets ready to take action. In the action phase interventions are followed aiming to fulfill proximal goals. After 6 months of actions the maintenance phase is reached, habit is created and should be maintained until termination stage, when user has no desire to return to old behaviors. Interventions related to different stages require different strategies and design. In this paper we will focus on discussing the development and design of digital interventions for action phase, i.e., the formation of a new habit aiming for getting to maintenance and hopefully termination stage (15). Interventions could be further supported by educational and motivational elements to enhance adherence. Further considerations for development of mHealth interventions can be found in (16).

1.1.2 Goal setting and habit formation

One promising way to lead individuals towards behavioral change and habit formation is through goal setting strategies. Distal goals (long-term) are achieved through a set of proximal (short-term) goals (17) that can develop into tiny habits (18). Distal goals are typically used for action planning and are formulated to enhance intrinsic motivation (19), whereas proximal goals are formulated more specifically to guide daily behavior. For proximal goal setting, the SMART (specific, measurable, achievable, relevant, and time-bound) framework can be used (20). A plan to achieve the distal goal thus consists of a hierarchical set of executable proximal goals that, when performed regularly, foster long-term lifestyle changes (21).

For example, in the context of a digital intervention aimed at improving physical activity, a distal goal could be stated as “*Be physically more active*,” while a corresponding SMART proximal goal could be: “*I will walk 7,000 steps at least 5 days in a week for 4 months, to increase my physical activity*.” This goal is specific about the action, measurable and time-bound. The relevance and achievability however, are subjective and must be tailored to the individual context and capabilities. When applying mHealth solutions in clinical practice, incremental sub-goals should be consulted with clinicians and aligned with the patient's medical profile. Additional aspects for designing effective SMART goals are addressed by (22).

Along with structural goal-setting, the psychological mechanisms behind motivation and habit formation should also be considered for digital health interventions design (23).

Formation of new healthy habits, simultaneous to breaking existing unhealthy ones (24), is facilitated through repeated behavior in stable contexts (25, 26). This is often described as cue-routine-reward loops (27). Digital health interventions can reinforce these loops by providing contextual cues and reinforcing feedback. Daily actions prompted by digital interventions can be complemented by reflective processes occurring over longer behavioral episodes, allowing for temporal alignment with real-world behaviors and capturing the dynamic interplay between behavioral and cognitive constructs (28). Moreover, theories of motivation, particularly the distinction between intrinsic and extrinsic motivation, play a pivotal role in long-term adherence. While extrinsic motivators such as reminders or rewards may initiate a behavior, the transition toward intrinsic motivation (e.g., walking because it feels good or aligns with one's values) is crucial for sustainability (29). Self-determination theory further suggests that behaviors are more likely to be maintained when they support a sense of autonomy, competence, and relatedness (30, 31). In mHealth applications, fostering *autonomy* can be facilitated through offering personalized goal setting and giving users choices in how they engage with interventions (e.g., selecting preferred types of activity or timing of reminders). *Competence* can be enhanced through clear feedback in form of e.g., progress tracking, with positive reinforcement features like daily streaks or milestone badges. To support *relatedness*, social features such as peer support groups, sharing achievements, or connecting with coaches and clinicians can help users feel part of a community. Additionally, adaptive interventions that evolve with the user's behavior and needs over time can further enhance engagement. By aligning design elements with psychological principles, mHealth tools can move beyond compliance to enable meaningful, sustained behavior change.

2 Methods

In this section, we outline key methodological concepts: the causal inference framework for unbiased effect estimation and multi-armed bandit algorithms for adaptive decision-making. These concepts provide the foundation for the design of intervention components and evaluation strategies discussed in the following sections.

2.1 Causal inference

Many questions in health and behavioral sciences are inherently causal rather than purely associative (32). The aim of causal inference is to move beyond mere associations between variables, towards understanding and quantifying the causal effect of an exposure on an outcome. It distinguishes correlation from causation by relying on formal assumptions, conceptual frameworks, and methods that allow for the estimation of causal effects even in the presence of confounding, selection bias, and measurement error (32). Moreover, unlike purely predictive models, causal methods also seek to answer *counterfactual*

questions: What would have happened to an individual or a population if a different action had been taken?

A foundational perspective in causal inference is the potential outcomes framework, also known as the Neyman-Rubin model (33). In this approach, each unit (e.g., an individual) is considered to have multiple potential outcomes, one for each possible treatment condition. For instance, a patient might have one outcome if treated with a drug and another if given a placebo. The causal effect is defined as a contrast between these potential outcomes, such as their difference or ratio. However, only one outcome is observed for each unit (the factual outcome), while the others remain unobserved (counterfactuals) (34). This constitutes the fundamental problem of causal inference: the impossibility of observing all potential outcomes for a single unit.

To make causal statements despite this limitation, researchers must rely on *assumptions* and designs that allow *identification* of causal effects. Central to this are the notions of *treatment* (or exposure), *outcome*, and *confounders* (variables that influence both treatment and outcome). Confounders can create spurious associations if not properly controlled, making their identification and *adjustment* essential for unbiased causal effect estimation (32).

Complementary to the potential outcomes approach is the use of graphical models representing causal relationships between variables. These graphs encode assumptions about the underlying data-generating process and enable formal reasoning about causation through graphical criteria such as d-separation and backdoor paths (32). When the structure of the graph is known or estimated, these tools provide a basis for identifying causal effects and designing appropriate adjustment strategies. A framework unifying the two paradigms, called single world intervention graphs (SWIGs), was presented in (35).

2.1.1 Structural causal model

A central development in causal inference is the formulation of structural causal models (SCMs). SCMs formalize causal assumptions through structural equations that describe how each variable is generated from its causes and an exogenous error term (36). The relationships among variables can be represented graphically as a directed acyclic graph (DAG), where nodes correspond to variables and directed edges represent direct causal influences. The DAG not only captures causal mechanisms but also encodes conditional independence relationships implied by the underlying causal structure. Understanding the graph structure is fundamental, as it dictates how we can adjust for confounding, predict the effects of interventions, and identify causal effects from data.

2.1.1.1 Causal discovery

In many applications, the true causal graph is unknown and must be inferred from observational or interventional data. Causal discovery, also known as structure learning, attempts to learn the structure of the underlying DAG, by estimating its equivalence class. Constraint-based methods, like the Peter-Clark (37) or Fast Causal Inference (38), infer graph structures by systematically testing conditional independencies, while score-based methods

search for the graph that best fits the data according to a predefined scoring criterion. Whenever possible, a background knowledge should be incorporated, to restrict the search space and enhance causal discovery algorithm. If the variables are collected over multiple time points, causal discovery can be guided by known time orderings (39). Optimization-based methods, such as NOTEARS (40), on the other hand, employ gradient-based techniques, but currently lack built-in support for incorporating background knowledge. An overview of causal discovery methods can be found in (41). The output of causal discovery provides a foundation for subsequent causal analysis and inference.

2.1.2 Effects and inference

With the potential outcomes and SCM frameworks in hand, we now focus on the task of estimating causal effects from the data. First, the DAG is used for identification of valid adjustment set via the back-door criterion, i.e., the smallest collection of covariates that blocks all noncausal paths from treatment to outcome (32). This step ensures that, under the assumptions of no unmeasured confounding and positivity, the target causal estimand (such as the average treatment effect) is identifiable (32).

Once the adjustment set is determined, the causal effect can be estimated through methods such as regression adjustment, where the outcome is modeled as a function of treatment and controls (confounders), and the predicted outcomes under different treatment levels are compared (42). The effect sizes of control variables, however, are unlikely to have causal interpretation themselves (43). Alternatively, inverse-probability weighting (IPW) (44) can be used, where each unit is weighted by the inverse of its probability of receiving the observed treatment, creating a pseudo-population in which treatment is independent of measured confounders. Another approach is the use of doubly robust methods (45–47), which combine outcome modeling with IPW to provide protection against misspecification in either of the models. Also meta-learners (48) utilizing machine learning techniques to estimate causal effect can be used. To estimate all effects on the graph, a (causal) Bayesian network (49) approach can be employed. In settings where the DAG is unavailable or incomplete, the intervention calculus when DAG is absent (IDA) method (50) can be used. This approach estimates the lower bound of total causal effects from observational data in high-dimensional settings, even without a fully specified causal graph. Additionally, by employing causal fusion, we can integrate data from both observational and experimental settings, as well as from multiple sources, enabling more robust causal inference across diverse datasets (51).

Finally, to assess the robustness of the derived findings to violations of unconfoundedness and to partially identify effects when point identification fails, sensitivity analysis can be performed and bounds estimated. The aim of sensitivity analysis is to quantify how strong an unmeasured confounder would have to be to overturn derived conclusions, using approaches such as Rosenbaum's bounds (52) or the E-value (53). This allows for reporting, for example, the minimum strength of association that an unobserved variable must have with both treatment and outcome to explain away the estimated

effect. When unmeasured confounding cannot be ruled out and point identification is questionable, bounds on the average causal effect can be estimated, e.g., via framework proposed in (54), and further advancements discussed in (55). Together, these sensitivity and bounding analyses complement the point estimates obtained under unconfoundedness and enhance the credibility of causal conclusions.

2.1.3 Longitudinal setting

Estimating individual treatment effect from longitudinal data, i.e., repeated measurements, can be done by utilizing described concepts. In digital health, an N-of-1 analysis of self-tracked data (56) can be employed to evaluate causal effects in longitudinal observational data, by utilizing the counterfactual framework. Findings from individual N-of-1 analyses can then be combined to provide population-level treatment effect estimates (57). The main assumptions for these methods include the consistency of exposure assignments over time, no interference between individuals, and a proper specification of the counterfactual model. Another approach for evaluation of causal effects in time series data is Granger causality (58), which can be extended to non-linear effect evaluation with copulas (59). Furthermore, if the structure is unknown, several data-driven approaches, utilizing variants of constraint and score-based causal discovery methods, were proposed (60–62).

2.2 Multi-armed bandit

The multi-armed bandit (MAB) problem is a framework from reinforcement learning (RL), designed for *sequential decision-making* under uncertainty (63). In each time step $t \geq 0$, an *agent* with K arms (options) selects one arm $k \in \{1, \dots, K\}$ according to a *policy*, and performs an action $a_t = k$. A *policy* is the agent's strategy for selecting actions based on the information available up to time t . It can be a deterministic or stochastic rule that maps the agent's current knowledge, typically the history of past actions and received rewards, into a choice of the next action. In the next time step $t + 1$, as a consequence of its action, the agent receives a numerical reward r_{t+1} , and updates the estimate of an expected reward for that action, which is assumed to be *stationary*. The agent's goal in MAB setting is to maximize expected total reward (63). Moreover, *regret*, defined as the difference between the reward that would have been obtained by always selecting the best arm and the reward actually accumulated, quantifies the cost of learning under uncertainty (64).

The true reward distributions associated with each arm are initially unknown and must be estimated through interaction with the environment (65). This introduces the exploration-exploitation trade-off: the agent must decide between *exploring* less certain arms to gather more information about their rewards and *exploiting* the arm that currently appears most promising. Several algorithms have been proposed to address the exploration-exploitation dilemma. Greedy algorithms prioritize the action with the highest estimated reward, favoring exploitation, whereas non-greedy strategies allow for occasional

exploration to prevent premature convergence to suboptimal arms (63). Different MAB algorithms balance exploration and exploitation in various ways, with each approach yielding distinct theoretical and empirical properties (66). The process is memoryless, meaning the agent's decisions at each time step depend only on the current policy, simplifying the analysis and implementation of learning algorithms.

2.2.1 Contextual and causal MAB

To improve the performance of MAB algorithms, contextual information can be incorporated both in the action selection process and within the learning phase, such as informing the reward design (67). In this scenario, the policy is influenced not only by the agent's previous actions but also by the *state* s_t of the environment, which is assumed to be stationary in MAB setting.

In RL, action selection is based on the current state and policy (63). By introducing structural assumptions about the environment through SCM and an associated causal diagram representing actions of agent, this task evolves into causal RL (68). This can be reduced to causal MAB (69), which can exploit assumed SCM for better action selection if all variables are manipulatable. Moreover, an extension to handle also cases when not all variables are manipulatable was introduced in (70). This approach enables evaluation of how a decision-maker should intervene to optimize non manipulatable outcome which is causally connected to manipulatable variables.

2.2.2 Dynamic environments

RL algorithms, including MABs, assume stationarity, i.e., the environment and reward distributions remain stable over time. This assumption however typically does not hold in dynamic environments, such as behavioral change programs. To tackle this, most recent observations can be treated as more important, by (possibly non linearly) weighting them more than past observations. For instance, this can be done with a sliding window approach, by setting the weight of past observations too far in the future to 0, effectively ignoring them, while only considering most recent ones.

A key challenge that remains open is the selection of an appropriate window size for adaptation, effectively a critical hyperparameter of the employed algorithm. The window size directly affects the trade-off between responsiveness to recent changes and robustness to noise. A small window can enable fast adaptation to sudden behavioral changes (e.g., illness or routine disruptions), but may result in instability due to overfitting to short-term fluctuations. Moreover, shorter window would lead to higher uncertainty (higher variance) in posterior distributions due to lower sample size, prioritizing exploration over exploitation in sampling. Conversely, a large window offers more stable estimates but may obscure important shifts in behavior. In practice, choosing the optimal window size depends on the expected rate of behavioral change and the variability of user engagement. A possible strategy to handle this could be fixing it to one plausible size according to previous studies or expertise. Another option would be selection of the window size in a data-

driven way by discounted history weighting, or by exploiting models that can utilize past observations to modify the window size in response to shift in reward distribution (71, 72). An empirical comparison of possible strategies should be done to better tailor digital health solutions to the target population. Strategy of the window size adaptation is an important aspect to consider, as it directly impacts the algorithm's ability to personalize interventions effectively over time. In the clinical and behavioral setting, window size choice could also take into consideration the need to reflect domain-specific rhythms, such as circadian cycles or weekly routines, to better capture meaningful patterns in individual's behavior.

2.2.3 Thompson sampling

Thompson sampling (TS) (73), a type of MAB, is a Bayesian approach to sequential decision-making under uncertainty, addressing the exploration—exploitation trade-off (72, 74). This algorithm operates by recomputing the posterior distribution after each received reward and thus building an implicit profile reflecting user's preferences. Thanks to these two properties it is a reasonable choice for applications in designing of digital health interventions.

In TS, reward likelihood for each arm is represented by an unknown probability distribution. After performing an action by the selected arm and observing the associated reward, the posterior distribution over this arm is updated. To facilitate easier posterior updates, it's advantageous to choose a conjugate prior distribution (75). If no information is available at the beginning, an uninformative prior should be chosen to start with, otherwise available data can be used to better initialize it. In case the reward of the k -th arm is binary, it can be described by a Bernoulli distribution with probability of success p_k , which can be modeled by Beta distribution $Be(\alpha_k, \beta_k)$ (76). In this case, for posterior update a Beta-Bernoulli is a convenient choice, as Beta distribution is a conjugate prior (75). For an uninformative prior, a Beta distribution with both α_k, β_k equal to 1, i.e., a uniform distribution on support $<0, 1>$, can be used. For the outcome of count data (e.g., Poisson) and with possibly high frequency of zeroes, count and zero-inflated models can be incorporated (77).

The action selection process in TS involves sampling from the posterior distributions over arms and selecting the arm with the highest sampled probability of success. If the environment is not stationary, a sliding-window TS (78) algorithm can be used. Posterior distributions are then computed only from the M most recent actions and the associated rewards observed, assuming stationarity in this window. The window size can be dynamically selected, as discussed in Section 2.2.2. Moreover, updating posterior estimates in TS in batches (79) instead of daily could help maintain more consistent estimates of reward distributions across arms, albeit at the expense of higher variance.

If distributions over arms for TS in each step depend also on the contextual variables, the contextual TS (80) can be used. Furthermore, relevant variables used for modeling user's state can be transformed to latent variables before applying it to TS (81), using statistical or machine learning models. If unobserved

confounders are present, an adaptation of TS proposed in (82) can be employed. Another variant of TS capable of handling online environments was proposed in (83). Ultimately, the choice of the specific TS variant, or other MAB, is task-dependent and should be guided by both technical feasibility and specific requirements of the application.

3 Personalized adaptive interventions

In mHealth applications, a personalized *nudge* is a brief persuasive intervention, that encourages and motivates users towards a specific behavior (84). Intervention aiming for behavioral change can be delivered in a text form through push notifications, in-app messages or through conversational interfaces. Also haptic feedback (85) can be used, by signaling e.g., milestones or providing real-time feedback as done in (86). Moreover, providing interim study results to participants (87) could enhance engagement. Other approaches not involving text itself that can serve as an intervention can be within a serious game, either via self-monitoring (88), rewards or other designs (89), or in other ways depending mainly on technical availabilities. The aim of intervention has to be clear, and motivational and persuasive techniques can be employed to enhance its efficacy. In this section, the opportunities and challenges in the design of digital interventions, based on the various data sources in mHealth applications, are discussed. To enhance the efficacy of digital interventions and user engagement, multi-armed bandits, introduced in Section 2.2, can be involved. MAB is preferred over RL thanks to its ability to adapt faster in the settings with little user interaction data in the early stages, prioritizing user retention through fast optimization. By focusing only on immediate feedback and avoiding the complexity of modeling long-term behavior, MAB allows for quicker adaptation and leads to lower regret compared to full RL.

3.1 Data sources

Collected data can be divided into three categories based on their source and characteristics: objective, subjective and contextual. Objective data are those that can be automatically collected through wearable devices such as smartwatches and smart rings (90). Examples of objective data collected from wearable devices are, just to name a few: steps, resting heart rate, heart rate variability, bedtime and wake up time, and physical activities. Objective data can also be collected through smartphones and mHealth app usage metrics (total time spent using an mHealth app, the way in which the app is used, ...). Subjective data are, on the other hand, collected directly from the user, either through questionnaires, conversational interfaces, or digital diaries. Subjective data can then be used to derive *explicit* contextual information and to better understand objective measures. Furthermore, *implicit* contextual information, such as weather and date-related features (day of the week, season, ...), are derived from smartphone background data (91). For our

example, data about physical activity can be obtained either as number of steps per day, or as activity logs (type, distance, other metrics related to the activity) or as intensity minutes (time spent in heart rate intervals). Daily steps will be used for evaluation of proximal goal, and other metrics will be relevant for evaluation of improvement in distal goal. Most wearables can provide this type of information, but it can be collected also by subjective measures, although with higher uncertainty.

3.1.1 Missing data

In the mHealth applications, missing and noisy data are a common problem. Non-availability of some data sources can be present, either due to users' lack of technology adoption, missing resources (wearable device not used or data collected only passively with wearables but not using other mHealth app components) or because of regulatory or technical restrictions (92, 93). The behavioral change system must be then designed to be scalable and usable for everyone even with a minimal data flow, while ensuring availability of relevant data sources. This structure entails a range of components within the system, from establishing minimal data for the basic version to leveraging all available data sources for more advanced features. Such design ensures scalability, accommodating varying degrees of user system adoption (94).

Even if the adoption of all features was successful, gaps and noise in longitudinal data will persist. These gaps can arise from factors such as non-compliance, where patients fail to adhere to wearable device usage or mobile app protocols, technical issues with the devices themselves or simply during charging the device. To deal with non-compliance, various approaches can be used, involving design considerations to increase adherence, e.g., with gamification or reward system (95, 96), or financial incentives (97). Further, to encourage adherence in digital diaries, notifications sent at random or scheduled times (98) or based on the contextual information about user's current state (99), if available, can be incorporated.

Handling missing data presents a critical challenge for the functioning of online behavioral change algorithms, particularly in digital health interventions. These "data holes" can disrupt real-time decision-making and compromise personalization and efficacy of interventions. Authors of (100) proposed a framework utilizing auxiliary variables to overcome this issue for handling missing accelerometer data in the outcome variables in trials, which could be adapted to other variables as needed. Further, in (101), expectation-maximization approach is compared to multiple imputation technique for missing step count data. To manage noisy data from wearables, signal processing approaches, such as filtering or denoising can be also utilized. If historical data about the user is available, imputation by utilizing similar context from previous days could be an option, however missingness mechanism might not be at random and this could lead to incorrect conclusions. Importantly, any behavioral decision based on incomplete data may be biased, and prolonged missingness can render the system non-functional. Therefore, ensuring data continuity, by the combination of technical means

and fostering user adherence, remains a key operational priority in the design of robust digital interventions.

For the post-analysis, i.e., evaluation of distal goals, and when deriving clinical trial endpoints, statistical approaches for missing data including within-patient imputations across common time periods, functional data analysis, and deep learning methods can be used, as presented in (102). For longitudinal studies with dropouts or truncation, the approach utilizing inverse-probability weighting and generalized estimating equations proposed by (103) can be considered. Other considerations for prevention and treatment of missing data in clinical trials are summarized in (104).

3.2 Just-in-time adaptive interventions

The most important aspect of digital health interventions is *when* they are delivered, to aim for the optimal outcome, i.e., behavioral response. Multiple *decision points* can be thus defined, either during the day, week or other relevant cycle, that is targeted through the definition of proximal goals. These do not have to be fixed, subject to technical availabilities they can be also event-based, if online data are available. That is however unrealistic in most of the mHealth applications, and this work will focus on the selection from the set of decision points. The simplest timing modality is at a reasonable fixed time, or a random selection from multiple options. This is especially useful for new users, when no relevant information is available for the decision. If available, contextual information should be used to better decide if to deliver the intervention at each decision point.

To increase efficacy of the intervention, minimize disruption to daily life and avoid frustration or disengagement, user's *internal state* should be determined prior to sending it (105). Nahum-Shani et al. (105) defines states of *vulnerability/opportunity* that aims to find a period of susceptibility to negative/positive health behavior changes, and *receptivity* targeting individual's willingness to receive the support delivered with the intervention. These concepts are explained in more detail also in (99). For instance, a reminder to walk more, as in our example, while driving should be avoided, as attempted in (106). Such intervention would be not only disengaging, but could also have a negative impact on the trustworthiness of the system's decisions.

The Just-In-Time Adaptive Interventions (JITAI) framework, introduced by (99), addresses this challenge by tailoring intervention timing to a user's current internal and contextual state. The decision rules are driven by defined thresholds for these *tailoring variables* and determine not only whether to intervene, but also the *type of intervention* to be delivered (105). Once the intervention is delivered, a proximal outcome (e.g., behavioral response) is observed. It is crucial that the timing of this outcome measurement aligns with the distal goal of the intervention, such as sustained behavior change (107). A meta-analytical review (108) of studies where JITAI were used concluded efficacy of this solution.

To further enhance efficacy of interventions, a (contextual/causal) MAB can be used to inform decisions in JITAI, as discussed in (109, 110), while keeping one arm as “no

intervention.” By utilizing TS, which addresses exploration—exploitation trade-off, user engagement can be enhanced by targeting habits with exploitation, while also new behaviors would be hopefully promoted via exploration. Moreover, a possible approach to mitigate user burden while sending interventions at various time points involves taking into account the cost of each intervention. This can be achieved by integrating a surrogate reward function into a MAB framework, as suggested by (111). The design of the reward function, which guides the selection of delivered interventions over time, must be tailored to the application, where this algorithm should be used.

3.3 Behavioral profiling and lifestyle quantification

To further enhance the behavioral change through personalized interventions, behavioral profiles created from available data can be incorporated (112) and used as a context and to describe user's internal state. Behavioral profiles are collections of user preferences, objective and subjective information, contextual features, and the statistical and causal relationships between them, both at the individual and population level. These can be further summarized into behavioral patterns describing individual's current habits. Formed behavioral profiles can then be used not only to inform content of personalized intervention, but also to improve timing, frequency and triggering conditions for interventions, i.e., to inform JITAI (107). The selection of variables and relationships relevant for behavioral profiling is based on the application's purpose, available data, and technical, regulatory and ethical restrictions.

Simplest behavioral profiles are explicit profiles, computed directly from the data as a vector of proportions of selected behavior, as proposed in (84). If a relevant information is available in the user's behavioral profile, timing can be even better personalized towards usual habits. For instance, a suggestion to use the bicycle instead of driving car for commuting to work, sent right after the usual wake up in the morning, could be particularly effective. Conversely, recommending that someone go for a walk at their most likely time, when they would already do so, doesn't foster the formation of a new habit. Additional value can then be provided, such as suggesting a new, longer, or more challenging route on the way back home to boost physical activity levels.

In order to better understand complex patterns and interactions in the collected data, statistical and causal modeling tools for repeated measurements, as introduced in Section 2.1.3, can be employed. As the data are collected longitudinally, N-of-1 design can be used to understand effects in longitudinal self-tracked data (56) but also to understand individual effect of previously sent interventions. Other machine learning techniques, such as generalized mixed effect regression, tree-based methods, neural networks or XGBoost can also be used for pattern recognition (113). Furthermore, a graphical representation of lifestyle, revealing structure among variables, can be estimated. To learn the structure of a causal graphical representation, causal

discovery methods, introduced in Section 2.1.1.1, can be used, as done in a data-driven way e.g., in (114). The effects and bounds are then estimated based on the derived structure. Learning effect from causal graphical structure is particularly useful in cases where intervening directly on the variable associated with the desired behavior is not feasible. Intervening on other variable which has direct causal effect on it, should consequently lead to a change in the desired direction for the target variable (115) as well. Moreover, by utilizing counterfactual framework on learned causal graph can be used for simulating of interventions and their consequences (116).

3.3.1 Health trajectories

Health trajectories can be identified using multi-modal information (117). These can be used not only for tracking disease progression, but also for assessing behavioral changes and informing the size of sliding window in MAB, and evaluating of effects of observed variables in a long-run. The identified personal patterns can then be used for improving personalisation in the interventions and sub-goal hierarchy design. Further, by identifying similar users based on their trajectories, predictions can be utilized to adapt the behavioral change program design to increase chance of effectiveness with growth mixture models (118). In future research, such estimated dynamic patterns could be incorporated to the post-analysis inference by including time-varying effects.

3.4 Goals, progress and feedback

The distal goal is divided into an ordered sequence of smaller executable proximal goals, as addressed in Section 1.1.2. A hierarchical set of proximal goals directs the trajectory of incremental daily lifestyle improvements, leading to the establishment of a new habit. In our example to improve physical activity, the proximal goal in SMART formulation can be “*I will walk 7,000 steps at least 5 days in a week for 4 months, to increase my physical activity.*” Based on the currently available data in the system, a progress towards this proximal goal is evaluated by fulfillment at each decision point. The computed progress can be then used as a *feedback* component in the intervention. Feedback contains an objective information encoding either progress towards the goal in form of an absolute or relative value, or an information about fulfilling the goal, strike of achieving goals, or any other option that can enhance motivation. Also comparison to similar people or one’s history could be incorporated (119). The feedback is thus one of the motivational components and should be selected in a personalized way to enhance efficacy of the intervention.

Furthermore, a proximal goal can be regularly re-evaluated, and adapted to be either more challenging or easier. This depends on whether it was fulfilled for a defined number of days in the last period, but also on the target population for behavioral change system, and in clinical settings should always be guided mainly by medical professionals. In any case, proximal goals should be updated and aim to lead to distal goal. When the

distal goal is reached, a maintenance behavioral change stage begins.

3.5 Motivational and persuasive component

Timing, factual content and type of intervention delivery were already discussed. In case of visual intervention, another design considerations can be made. Especially if the intervention is delivered in a form of text, i.e., via notification or through chatbot, motivational and persuasive components can be incorporated to enhance user’s engagement (120).

Firstly, the way how the message is presented can be used as a motivational component. In (84) this is facilitated through so-called triggering conditions, which is a set of rules that depend on the evaluated progress towards measurable proximal goal. This can be for example a *praise*, if user already fulfilled the proximal goal, or some type of *suggestion* if the proximal goal was not fulfilled yet. The rules for triggers must be compatible with the aim of the system and should be designed by professionals in the field. Triggering conditions serve not only as motivational component, but also give a base for evaluating whether the user’s behavior is aligned with the aim of the intervention, as described in the example depicted in Figure 1, and can thus help formulate the reward for multi-armed bandit. An example of triggering logic, including triggering conditions and the timing for assessing intervention effectiveness related to our example, is presented in Table 1. Conditions should be always designed based on the available data and target population, and should take into account also contextual state of the user. In special cases with too little available data, this could reduce to a simple scenario with only 1 row.

The motivational component of the intervention can include real-time and historical observations, whereas a persuasive strategy does not need to rely on the collected data, as it can be solely a persuasive statement. A set of persuasive strategies needs to be defined a-priori. For example Cialdini’s 6 persuasive strategies (121) can be used, as done in (122). A selection of persuasive strategies to enhance adherence to the lifestyle change program can be chosen according to the application purpose and target population. The strategy for the persuasive component, guiding the formulation of the intervention to be sent, can be selected in different ways. A straightforward approach could involve random selection or predetermined fixed probabilities. To allow personalization, a MAB algorithm can be employed. The application of TS for personalized selection of persuasive strategies has been shown to enhance adherence (122), and was also successfully applied to select one of the 6 Cialdini’s principles (123) as a way of *implicit* personalization. This technique is chosen because of its proven efficiency for exploration-exploitation trade-off, and because the posterior distributions recomputed in each step can be viewed as implicit persuasive profiles. Despite the persuasive profiles being only approximate and not correct estimates, they are sufficient for the purpose of the automatic decision-making for interventions (123). For a new user, prior distributions can be either

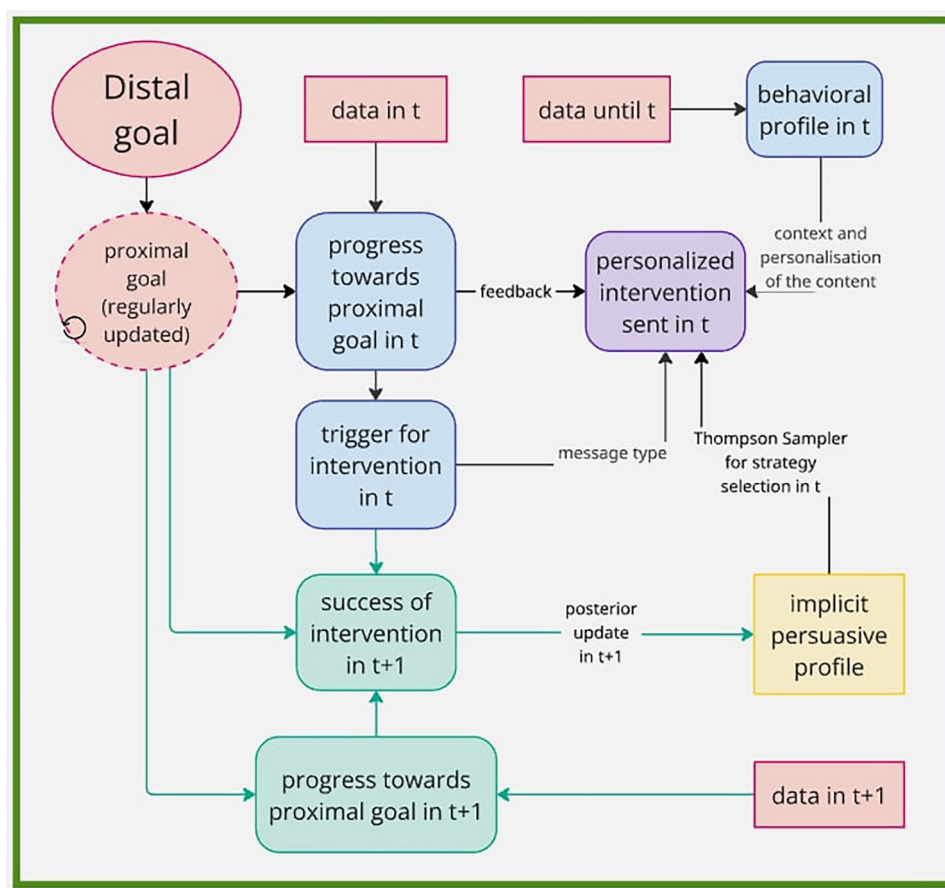


FIGURE 1

Example of intervention framework using MAB for selection of persuasive strategy in text message. Information from the system are in red, decisions made for intervention creation in time (t) are in blue. Design for update of the implicit profile in time (t+1) is in green. Created using <https://miro.com/>.

TABLE 1 Table of example triggers. Based on the condition, given by progress towards proximal goal, the type of the trigger is selected. The proximal outcome is then evaluated at corresponding time.

ID	Type	Condition	Evaluated when	Example
1	Praise	Progress $\geq 100\%$	End of the next day	Congratulations!
2	Nudge, encouragement	Progress $\in [80\%, 100\%)$	End of the day	Almost there!
3	Suggestion	Progress $< 80\%$	End of the day	Take a walk.

uninformative, or can be jump-started by using prior knowledge. This can be achieved either with relevant questionnaire filled in by the user upon logging in to the system for the first time, as done in (122), or by comparing to similar participants previously enrolled to the study based on the relevant metrics.

Returning to our example related to physical activity, having fixed a daily goal of 7,000 steps at least five times a week, the nudging message could be triggered if the user has taken less than 7,000 steps today and/or has not yet met the step goal in at least five times this week. Furthermore, assuming we know, from contextual/demographic information gathered, that the weather is currently sunny and that the user has a dog, the message may be formulated as follows: *Today you already fulfilled 5,000 steps out of your daily step goal of 7,000 steps. The weather is nice, how*

about taking your dog for a walk? There are only a few hours left today to fulfill your goal! This message, depicted in purple in Figure 1, is composed of several components: feedback (5000 out of 7,000 steps), contextual information collected automatically from location data (weather is nice), contextual information explicitly asked to the user in previous interaction (take your dog for a walk) and a persuasive component (Only a few hours left today) selected by MAB.

3.6 MAB and reward function in mHealth

In the context of mHealth interventions, the integration of MAB algorithms requires careful specification of the reward

function. The reward reflects the observed outcome of an intervention and is essential for guiding the posterior update. In adaptive interventions, rewards are typically based on proximal outcomes, which inform the algorithm's learning and adaptation during the study. However, distal outcomes should also be considered when defining the reward structure, as they ultimately determine the intervention's effectiveness. Proximal outcomes are evaluated in an online manner based on the triggering condition, i.e., the logic or rule that determines when and why an intervention is delivered. This is illustrated in green in Figure 1. Distal outcomes are assessed during the post-study analysis but should influence the initial reward design to ensure alignment between short-term optimization and long-term goals.

In MAB for mHealth, the reward assigned to an action is determined using data collected after the action is taken. For instance, if a *suggestion* message is sent, the intervention is considered successful if the user performed the desired behavior by the end of the day or up to few hours from the delivery of the intervention. Alternatively, for a *praise* message, success might be defined as the user exhibiting the desired behavior the next day without further prompts. An example of such triggering conditions and their corresponding success criteria are summarized in Table 1. Reward definitions should be tailored to the intervention type, desired behavior change, and data availability. If the reward is binary (e.g., success/failure), a Beta-Bernoulli model is used for the posterior update, otherwise different suitable model must be selected. For further guidance on the design of reward functions in mHealth bandit settings, see (124).

Some considerations for designing of reward function for TS are discussed in (111). Further, examples utilizing MAB in mobile health in various domains involve (77, 109, 125–130). These can also serve as examples for reward design and related considerations. Other examples of RL methods for the application to digital health involve e.g., (131) where authors used Bayesian Mixed Linear Model to approximate the reward, and (132) for multi-agent RL in digital health domain.

3.7 Multiple lifestyle pillars

Designing interventions for lifestyle improvements is a complex process, especially if more than one lifestyle pillar should be addressed in a single application. Addressing multiple pillars at the same time could lead to an excessive number of interventions within a short period, potentially burdening the user. Several strategies could be employed to strike a balance between minimizing potential burden and intervening sufficiently to achieve the desired effect. Pillars can be addressed either individually, one-by-one, or simultaneously, with the primary focus on one at a time. As before, automatic selection with MABs could be incorporated to enhance personalization. The design of reward for such decision-making system would be however probably the biggest challenge.

Furthermore, it is worth to consider, that what could lead to the highest reward in short term, might not be the best decision

in long-term, as the aim of the system is to change behavior and maintain healthy habits. The reward must be thus designed to reflect both proximal and distal goal at the same time. Moreover, even if interventions would not work as intended at the beginning, they might be effective for behavioral change in the long run. For this, incorporation of full RL with delayed rewards (63) in the later stages of the behavioral change program could enhance overall efficacy.

4 Evaluation and inference

To understand the efficacy of the behavioral change system, a comprehensive statistical evaluation is essential. Interventions delivered over longer period of time can be evaluated either via randomized controlled trial (RCT) by assigning participants into control vs. intervention group, or on an individual level, allowing also for evaluation of individual components. In this section, various possible approaches and associated challenges will be discussed, starting from defining the research question.

For proper communication, an estimand must be formulated, providing precise definition of outcomes of the study and other relevant factors (133). For this, a PICO (Population/Problem, Intervention, Comparison, and Outcome) framework, which plays a crucial role in evidence-based medicine (134), can be employed. PICO aims to guide study evaluation, serving to clearly define research questions and inform model specification. In the *population* component, inclusion and exclusion criteria are defined. These have to be inclusive, fair and ethical, but at the same time aiming to answer the research question of interest, while avoiding selection bias. The aim is to enroll a representative sample of the target population. *Intervention* specifies type of intervention delivered, such as tailored digital nudges delivered via a smartphone application, or more generally, usage of the behavioral change system. The choice of *comparator* then defines competing intervention or a standard care. For evaluation of digital interventions this can be no intervention, or not using the system at all, as also only self-monitoring can be an intervention (88). Finally, *outcome* quantifies the research question and guides model selection and analysis. In the example of physical activity this could be daily step count and its change throughout the study. On top of the primary outcome of interest, secondary outcomes can be included, targeting e.g., subgroup-specific effects or time-varying effects reflecting treatment effect heterogeneity over time. To enhance reporting quality of randomized trials, dedicated guidelines should be followed (135).

The study is typically framed as a superiority design aiming to show benefit of intervention against comparator. Digital interventions may however not always outperform existing treatments in absolute effect size, but can still provide meaningful value if they achieve similar outcomes with improved usability or accessibility. For this, a non-inferiority design can be more suitable (136). In the non-inferiority trial, an alternative hypothesis is formulated through a non-inferiority margin. The confidence interval of the treatment effect must then lie outside

of the inferiority region defined by the non-inferiority margin (137). The non-inferiority margin, determined in advance, is typically selected based on the results from a previous study, or by an expert, and should depend on the severity of the endpoint (138).

4.1 Study design

Study objectives defined by PICO are then defining the quantity of interest. Typically this would be an average treatment effect (ATE), i.e., the causal effect of treatment on the outcome. To evaluate ATE on a population level, an RCT is typically used. RCT is a gold standard for estimating unbiased causal effects, as random assignment facilitates identification of the ATE under minimal assumptions (32, 136).

If the sample size required for a reasonable power of the test can't be met or ethical considerations do not allow for running an RCT, an observational study designed as one-arm feasibility study (139) can be performed to test preliminary hypotheses of associations and collect feedback. If longitudinal data were collected, an effect of interventions on an individual level can be evaluated, and changes in trends should be visible. Furthermore, potential research designs to evaluate efficacy and effectiveness of mHealth interventions are summarized in (140). The effect of a new treatment on the selected study endpoints collected through one-arm experiment can be then either compared to outcomes in the synthetic control arm (141) or to the accepted conventional value of outcome by a non-inferiority study design (138).

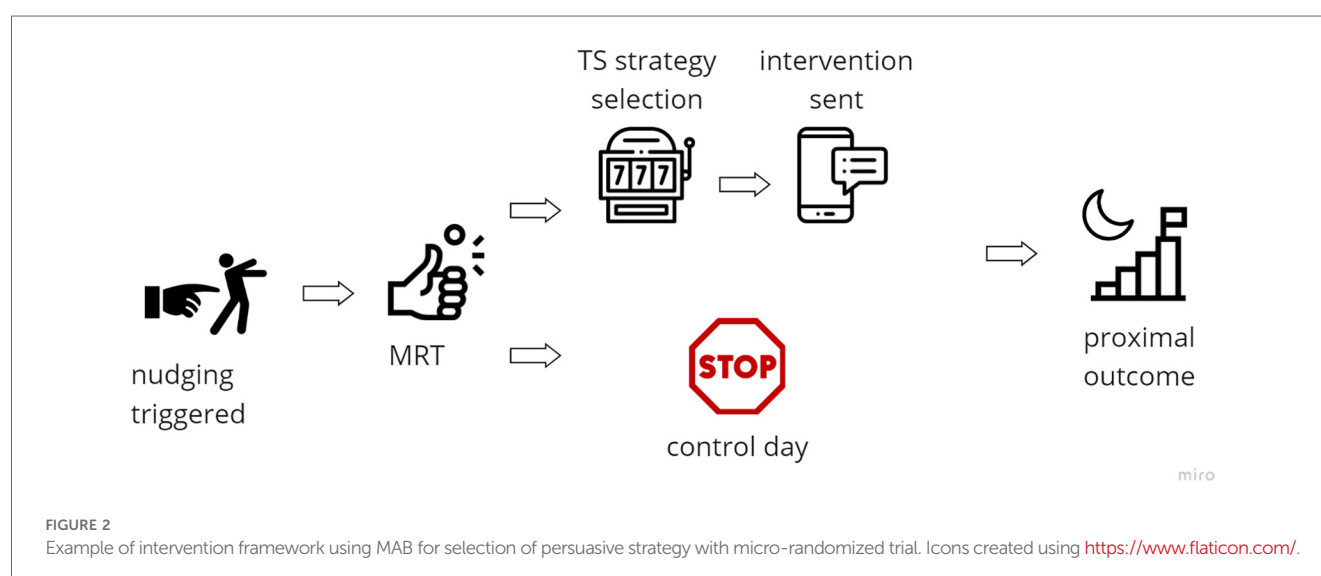
4.1.1 Micro-randomized trials

Post-hoc analysis in RCT provides an evaluation of efficacy of the overall behavioral change program and hypothesis testing, but the information about contribution of intervention components and information for future improvements is limited (142, 143). To understand effectiveness of interventions, a multilevel (fractional) factorial experiment should be employed (143).

To further evaluate the effectiveness of digital interventions at the individual level and in real-time, *micro-randomized trial* (MRT) offers a powerful experimental design that complements traditional RCT. While RCT assesses the ATE on distal outcomes across a population, MRT focuses on proximal outcomes, i.e., short-term responses to interventions delivered at specific decision points (144). MRT can decrease the sample size compared to a full factorial design, while still allowing for causal inference of intervention effects (144). In the MRT, participants are repeatedly randomized, often several times per day, to either receive or not receive an intervention at predefined decision points, using fixed probabilities. This repeated randomization allows for the estimation of time-varying and context-dependent causal effects at the individual level (144). An example of MRT design with persuasive component selected by MAB is illustrated in Figure 2.

MRT is particularly well-suited for informing JITAIs, where decisions about intervention delivery incorporate contextual features such as location, mood, or previous engagement (105). MRT supports various outcome types, including binary (145), continuous, and zero-inflated count outcomes (146), and facilitate analysis of time-varying moderator effects (147). MRTs have been successfully applied in domains such as physical activity promotion (125), alcohol reduction through timely notifications (148), and digital cardiac rehabilitation (149, 150). Considerations for sample size for MRTs are discussed in (151).

Although MRTs are designed for short term decision-making and personalization, they do not replace RCTs. Instead, both designs can be used in a complementary manner, for instance assessing proximal effects via MRTs and distal outcomes via RCTs in the same study (150). While MRTs are designed to assess the short-term effects of interventions, understanding their cumulative impact on long-term behavioral or health outcomes requires integrating longitudinal modeling with causal inference frameworks, including those that estimate individual treatment effects over time (99).



4.1.2 Adaptive designs

Traditional RCTs typically use fixed randomization ratios to assign participants to treatment or control groups. However, *response-adaptive randomization* (RAR) improves upon this by modifying allocation probabilities during the trial based on accumulated outcome data. The core idea is to assign more participants to treatment arms showing superior performance, thereby potentially enhancing patient benefit without sacrificing statistical validity (152, 153). RAR is particularly relevant in clinical trials where ethical considerations favor minimizing exposure to inferior treatments. To dynamically select allocations to arms, MAB approach can be used (154), aiming to maximize expected rewards (e.g., engagement, behavioral change). To enhance balance of exploration and exploitation in settings with satisfactory sample size, TS introduced in Section 2.2.3 can be used. For inference from data collected under MAB, incorporating an off-policy learning paradigm (155) could be beneficial. However, further research is needed to fully explore and refine this framework.

The concept of adaptive designs can be applied also to high-frequency interventions such as MRTs, where treatment options are randomized at multiple decision points over time. Instead of using fixed randomization probabilities, data from prior interventions can inform future allocations, optimizing intervention delivery based on individual participant responses. Further adaptation can be achieved through *Sequential Multiple Assignment Randomized Trials* (SMART), where participants are re-randomized based on earlier responses (156). SMART design combines optimization by JITAI, and utilizes MRT to also allow inference. Moreover, the behavior change program could be then formed by multiple components, focusing on different lifestyle pillars in parallel. The selection of the experimental design must be tied to design of interventions' components and desired inference; a more detailed overview can be found in (157).

In digital health, these adaptive designs can be embedded within a *Multiphase Optimization Strategy* (MOST), where interventions are first optimized for individual needs (e.g., through MRT or SMART) and then rigorously tested via RCT to ensure both individual tailoring and robust population-level evidence (142).

4.2 Statistical analysis and inference

A clear specification of estimand to evaluate intervention effect is critical for valid statistical analysis, especially in adaptive designs where both exposure and adherence can vary over time. *Intention-to-Treat* (ITT) assesses the effect of the assigned intervention, regardless of whether it was received, preserving the benefits of randomization and minimizing bias (158). On the other hand, *Per-Protocol* (PP) evaluates the effect among participants who fully adhered to the protocol (159). While potentially more relevant for assessing efficacy, this approach risks introducing selection bias due to non-random adherence.

To mitigate this bias, *instrumental variable* (IV) methods can be employed (160, 161). In randomized trials, the random

assignment itself serves as a valid instrument to estimate the causal effect of treatment receipt, also known as the complier average causal effect (CACE) (162). In digital interventions, this approach is particularly relevant when user interactions (e.g., notification clicks or app usage logs) can be used as proxies for actual treatment exposure (163).

Given that digital health studies typically involve repeated measurements over time, appropriate longitudinal models are essential. *Mixed-effect models* account for within-subject correlations and allow for estimation of both fixed effects (e.g., treatment) and random effects (e.g., individual-level variation) (164, 165). Depending on the outcome distribution, linear or generalized models should be used. These models can be estimated using methods such as restricted maximum likelihood (REML), the EM algorithm (166), or Bayesian inference (167). It is important to note that fixed effects can only be interpreted marginally if treatment assignment is independent of random effects (168). In addition, *survival analysis* methods such as Cox proportional hazards models can be used to evaluate time-to-event outcomes, such as time to dropout (169). Moreover, several statistical frameworks have been proposed to analyze MRT data, including weighted and centered least-squares estimators (170) and robust inference methods for longitudinal binary outcomes (145). To evaluate individualized treatment effects, facilitating personalized decision-making within the system, machine learning models and concepts from causal inference can be employed (171). Furthermore, *multilevel models* can estimate time-varying treatment effects of proximal interventions delivered repeatedly throughout the study under MRT design (144). Moreover, in the studies with MRTs, although not desired, a carry-over effect might occur and should be controlled for in the analysis. These can involve not only positive effects of interventions from previous days, but also experience of burden.

To ensure the validity of results, especially in the presence of multiple testing and repeated measures, all components of the statistical analysis, including model specifications and hypotheses, should be pre-specified before data collection. This guards against inflated Type I error (172). Furthermore, when evaluating relationships for which randomization was not applied, *causal inference tools*, introduced in Section 2.1.2, can be used to estimate effects from observational data. Approaches such as propensity score adjustment, IPW, or g-methods help reduce confounding. Moreover, findings from both randomized and observational components can be combined using a *causal fusion* framework (51), enabling the integration of multiple evidence sources to strengthen causal conclusions.

5 Discussion

This paper reviewed personalized adaptive interventions through the lens of MAB algorithms, focusing on their application in behavioral change contexts. MABs were selected due to their lower initial regret compared to full RL approaches, making them suitable for early-stage interventions where user engagement is fragile. Moreover, experimental designs allowing

for proper statistical inference were discussed, ranging from RCT and MRT for individual treatment effect estimation, up to adaptive designs. These methods provide a toolkit for balancing personalization and scientific rigor in digital health. The discussed framework and similar adaptive approaches pave a way for opportunities in personalization of mHealth applications not only for behavioral change, but also for chronic disease prevention and management.

Designing effective and scalable adaptive systems presents several challenges. Interventions must be engaging, ethically sound, and technically feasible. A key difficulty lies in delivering the right support at the right moment. Accurate user profiling can enhance the timing and personalization of interventions, yet further research is needed to fully leverage MABs in this area to maximize adherence and impact without compromising inference. Personalization, while essential for effectiveness, inherently limits generalizability to broader populations. Additionally, behavioral systems must be dynamic, continuously adapting to evolving user profiles and behaviors. This motivates techniques such as sliding windows and non-stationary models to ensure continued relevance and effectiveness. User retention is another critical factor. MABs can help reduce dropout in early stages, where it tends to be highest, by adapting quickly to user responses even with little data.

Despite the promise of goal-setting and motivational strategies, long-term engagement remains a significant challenge. Users may disengage due to a lack of perceived progress, goal misalignment, interface complexity, or poor fit with changing personal circumstances and life contexts (173–176). To address this, interventions should incorporate adaptive feedback loops and re-engagement mechanisms such as timely nudges and dynamic goal adjustments (177). Periodic reassessment of preferences and behavioral stages can help ensure the intervention remains relevant and supportive. While psychological perspectives are increasingly integrated into digital intervention frameworks (23, 28), further empirical research and evaluation across diverse domains of lifestyle change would strengthen the validity of these approaches. Gamification, socialization, education, and engaging feedback mechanisms may also foster deeper involvement. Additionally, fostering intrinsic motivation and self-determination in behaviors might be more effective than relying on extrinsic rewards and external regulation to drive behavior. Moreover, from a psychological standpoint, sustainable behavior change is better facilitated through the pursuit and attainment of learning and process goals, rather than focusing solely on the outcome goals (29). These aspects have to be however carefully translated into the design of AI-enhanced behavioral change systems.

Ethical and regulatory considerations are essential throughout the system's lifecycle. Interventions must ensure user safety, transparency, and autonomy. Users should retain control over their data, as in EU under the law of General Data Protection Regulation (GDPR, Regulation (EU) 2016/679), and be informed about how decisions are made. Systems must avoid unintended harms, such as delivering inappropriate recommendations or reinforcing harmful behaviors (178). Equity, inclusiveness, and

fairness are critical to building trust and ensuring broad societal benefit. Further ethical considerations for usage of automatic decision-making systems in healthcare are summarized in (179). For health-related applications, also regulatory compliance is paramount. Interventions should be developed in consultation with domain experts and follow relevant legal guidelines. In the U.S., the FDA provides regulatory guidance for digital health technologies (180), while the European Union's Artificial Intelligence Act (EU AI Act, Regulation (EU) 2024/1689) introduces specific requirements for AI-based systems. Notably, behavior change systems, though generally prohibited, are allowed in the medical domain under strict conditions. Certification under the EU Medical Device Regulation (MDR, Regulation (EU) 2017/745) may be required, depending on the intervention's intended use and decision-making autonomy. For example, automated goal setting in cardiac rehabilitation programs may require oversight from medical professionals to ensure safety. The use of generative AI, such as large language models (LLMs), also raises regulatory challenges. Automatically generated motivational messages may face restrictions, and should ideally rely on professionally vetted content. As discussed in recent literature, the growing role of LLMs in healthcare applications demands further research into appropriate safeguards and governance mechanisms (181). Evaluation of mHealth solutions also requires clear objectives, validated outcome metrics, and longitudinal monitoring. In digital interventions, as in other healthcare applications domains, RCTs are gold standard for evaluation of the system. When dealing with repeated exposures, MRTs offer advantages over traditional RCTs, such as smaller sample size requirements and the ability to assess individual intervention components. This allows for iterative refinement based on real-world data and user feedback, while adhering to the framework of RCT. However, further research is needed to understand generalizability of the results from MRTs. Additionally, the so-called Hawthorne effect (182), where participants alter behavior simply because they are being studied, can inflate perceived intervention efficacy. Therefore, future studies should include post-deployment evaluations in naturalistic settings to gauge real-world effectiveness.

This review highlights a rich and evolving landscape of methods for designing and evaluating adaptive digital health interventions. The main opportunity lies in the real-time personalization offered by MABs and related adaptive methods, which can improve engagement and efficacy while preserving analytical tractability. These approaches enable rapid learning from user behavior and offer fine-grained control over intervention delivery. However, this flexibility also introduces trade-offs. Personalization can reduce generalizability; adaptivity can complicate causal inference; and automated decisions raise ethical and regulatory concerns. While MABs reduce regret in the short term, they may not fully capture long-term outcomes or delayed effects, for which more complex RL methods might be required. Furthermore, deriving causal effects especially from observational data for constructing of behavioral models is a challenge, as it is not possible to derive causal effects purely from observational data without posing (sometimes untestable)

assumptions. Current tools for causal discovery offer promising results when the background knowledge is incorporated. However, incorrect direction of edge in the output can introduce incorrect assumptions for the subsequent identification and inference (183). While proper selection of confounders reduces bias in the causal estimates, selection of inappropriate controls, such as colliders or mediators, can result in estimates that are biased even more than uncontrolled estimates (184).

A promising direction for future research is the creation of personalized *digital twins*—virtual representations of individuals built from continuous wearable data and contextual information to capture behavioral trajectories (185, 186). Building of digital twins unifies the key ideas discussed in this review, including MABs, adaptive learning, causal inference, and interventions simulation, within a single, coherent framework. By simulating different intervention scenarios, digital twins can help optimize the timing and content of support while avoiding real-world risks. Integrating causal inference and adaptive algorithms into these models would allow researchers to test and refine intervention strategies virtually before deploying them in practice. This approach could enable highly personalized treatment plans, enhance safety, and accelerate development cycles. However, important challenges remain, such as ensuring model accuracy, managing computational complexity, protecting user privacy, and validating predictions against real-world behaviors. Future work should focus on scalable methods to build and update digital twins in real time, evaluate their advantages over conventional adaptive designs, and develop metrics to measure their precision and impact in digital health interventions.

In conclusion, adaptive mHealth interventions leveraging MAB algorithms, causal inference frameworks, and emerging digital twin technologies represent a promising frontier for personalized behavior change. Their success, however, depends on ongoing adaptation, rigorous evaluation, psychological insight, interdisciplinary collaboration, and compliance with regulatory and ethical standards. Future work should validate long-term efficacy, investigate generalizability across populations, and

explore broader applications across behavioral, clinical, and public health domains.

Author contributions

RŠ: Conceptualization, Investigation, Visualization, Resources, Formal analysis, Writing – original draft, Methodology, Writing – review & editing. ADR: Writing – review & editing. DM: Writing – review & editing. AT: Writing – review & editing. FDF: Supervision, Writing – review & editing, Funding acquisition.

Funding

The author(s) declare that financial support was received for the research and/or publication of this article. The present research was partially funded by the European Union and the Swiss SERI through the Eurostars funding project CUOREMA [grant number E115067].

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

1. World Health Organization. Noncommunicable diseases. *WHO Fact Sheet*. (2023).
2. Aboyans V, Causes of Death Collaborators. Global, regional, and national age-sex specific all-cause and cause-specific mortality for 240 causes of death, 1990–2013: a systematic analysis for the global burden of disease study 2013. *Lancet*. (2015) 385:117–71. doi: 10.1016/S0140-6736(14)61682-2
3. Chatterjee A, Prinz A, Gerdes M, Martinez S. Digital interventions on healthy lifestyle management: systematic review. *J Med Internet Res*. (2021) 23:e26931. doi: 10.2196/26931
4. Kumar S, Nilsen W, Pavel M, Srivastava M. Mobile health: revolutionizing healthcare through transdisciplinary research. *Computer*. (2012) 46:28–35. doi: 10.1109/MC.2012.392
5. Jusoh S. A survey on trend, opportunities and challenges of mHealth apps. *Int J Interact Mob Technol*. (2017) 11:73–85. doi: 10.3991/ijim.v11i6.7265
6. Wilhide III CC, Peeples MM, Kouyaté RCA. Evidence-based mhealth chronic disease mobile app intervention design: development of a framework. *JMIR Res Protoc*. (2016) 5:e4838. doi: 10.2196/resprot.4838
7. Menictas M, Rabbi M, Klasnja P, Murphy S. Artificial intelligence decision-making in mobile health. *Biochem (Lond)*. (2019) 41:20–4. doi: 10.1042/BIO04105020
8. Sannino G, Bouguila N, De Pietro G, Celesti A. Artificial intelligence for mobile health data analysis and processing. *Mobile Inform Syst*. (2019) 2019:2673463. doi: 10.1155/2019/2673463
9. Thomas Craig KJ, Morgan LC, Chen C-H, Michie S, Fusco N, Snowdon JL, et al. Systematic review of context-aware digital behavior change interventions to improve health. *Transl Behav Med*. (2021) 11:1037–48. doi: 10.1093/tbm/ibaa099
10. Kankanhalli A, Xia Q, Ai P, Zhao X. Understanding personalization for health behavior change applications: a review and future directions. *AIS Trans Hum Comput Interact*. (2021) 13:316–49. doi: 10.17705/1thci.00152
11. Rippe JM. Lifestyle medicine: the health promoting power of daily habits and practices. *Am J Lifestyle Med*. (2018) 12:499–512. doi: 10.1177/1559827618785554
12. Fallows ES. Lifestyle medicine: a cultural shift in medicine that can drive integration of care. *Future Healthc J*. (2023) 10:226. doi: 10.7861/fhj.2023-0094
13. Bandura A. Self-efficacy: toward a unifying theory of behavioral change. *Psychol Rev*. (1977) 84:191. doi: 10.1037/0033-295X.84.2.191
14. Prochaska JO, Velicer WF. The transtheoretical model of health behavior change. *Am J Health Promot*. (1997) 12:38–48. doi: 10.4278/0890-1171-12.1.38

15. Gardner B, Lally P, Wardle J. Making health habitual: the psychology of “habit-formation” and general practice. *Br J Gen Pract.* (2012) 62:664–6. doi: 10.3399/bjgp12X659466
16. Milne-Ives M, Lam C, De Cock C, Van Velthoven MH, Meinert E. Mobile apps for health behavior change in physical activity, diet, drug and alcohol use, and mental health: systematic review. *JMIR mHealth uHealth.* (2020) 8:e17046. doi: 10.2196/17046
17. Bandura A, Simon KM. The role of proximal intentions in self-regulation of refractory behavior. *Cognit Ther Res.* (1977) 1:177–93. doi: 10.1007/BF01186792
18. Fogg BJ. *Tiny Habits: The Small Changes that Change Everything.* Boston, NY: Harvest (2020).
19. Bailey RR. Goal setting and action planning for health behavior change. *Am J Lifestyle Med.* (2019) 13:615–8. doi: 10.1177/1559827617729634
20. Bovend'Eerd TJ, Botell RE, Wade DT. Writing smart rehabilitation goals and achieving goal attainment scaling: a practical guide. *Clin Rehabil.* (2009) 23:352–61. doi: 10.1177/0269215508101741
21. Shilts MK, Horowitz M, Townsend MS. Goal setting as a strategy for dietary and physical activity behavior change: a review of the literature. *Am J Health Promot.* (2004) 19:81–93. doi: 10.4278/0890-1171-19.2.81
22. Bowman J, Mogensen L, Marsland E, Lannin N. The development, content validity and inter-rater reliability of the smart-goal evaluation method: a standardised method for evaluating clinical goals. *Aust Occup Ther J.* (2015) 62:420–7. doi: 10.1111/1440-1630.12218
23. Pinder C, Vermeulen J, Cowan BR, Beale R. Digital behaviour change interventions to break and form habits. *ACM Trans Comput Hum Interact.* (2018) 25:1–66. doi: 10.1145/3196830
24. Wood W, Neal DT. Healthy through habit: interventions for initiating & maintaining health behavior change. *Behav Sci Policy.* (2016) 2:71–83. doi: 10.1177/237946151600200109
25. Lally P, Gardner B. Promoting habit formation. *Health Psychol Rev.* (2013) 7: S137–58. doi: 10.1080/17437199.2011.603640
26. Wood W, Rünger D. Psychology of habit. *Annu Rev Psychol.* (2016) 67:289–314. doi: 10.1146/annurev-psych-122414-033417
27. Duhigg C. *The Power of Habit: Why We Do What We Do in Life and Business.* New York, NY: Random House (2012). Vol. 34.
28. Zhang C, Lakens D, IJsselstein WA. Theory integration for lifestyle behavior change in the digital age: an adaptive decision-making framework. *J Med Internet Res.* (2021) 23:e17127. doi: 10.2196/17127
29. Ryan RM, Deci EL. Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *Am Psychol.* (2000) 55:68. doi: 10.1037/0003-066X.55.1.68
30. Ntoumanis N, Ng JY, Prestwich A, Quested E, Hancox JE, Thøgersen-Ntoumani C, et al. A meta-analysis of self-determination theory-informed intervention studies in the health domain: effects on motivation, health behavior, physical, and psychological health. *Health Psychol Rev.* (2021) 15:214–44. doi: 10.1080/17437199.2020.1718529
31. Ryan RM, Deci EL. *Self-Determination Theory: Basic Psychological Needs in Motivation, Development, and Wellness.* New York, NY: Guilford Publications (2017).
32. Pearl J. Causal inference in statistics: an overview (2009).
33. Rubin DB. Estimating causal effects of treatments in randomized and nonrandomized studies. *J Educ Psychol.* (1974) 66:688. doi: 10.1037/h0037350
34. Collins J, Hall N, Paul LA. *Causation and Counterfactuals.* Cambridge, MA: MIT Press (2004).
35. Richardson TS, Robins JM. Single world intervention graphs (SWIGs): a unification of the counterfactual and graphical approaches to causality. In: *Center for the Statistics and the Social Sciences, University of Washington Series. Working Paper.* (2013). Vol. 128. p. 2013.
36. Pearl J. *Models, Reasoning and Inference.* Cambridge, UK: Cambridge University Press (2000). Vol. 19. p. 3.
37. Spirtes P, Glymour C. An algorithm for fast recovery of sparse causal graphs. *Soc Sci Comput Rev.* (1991) 9:62–72. doi: 10.1177/089443939100900106
38. Spirtes PL, Meek C, Richardson TS. Causal inference in the presence of latent variables and selection bias. *arXiv [Preprint]. arXiv:1302.4983* (2013).
39. Petersen AH, Osler M, Ekström CT. Data-driven model building for life-course epidemiology. *Am J Epidemiol.* (2021) 190:1898–907. doi: 10.1093/aje/kwab087
40. Zheng X, Aragam B, Ravikumar PK, Xing EP. Dags with no tears: continuous optimization for structure learning. *Adv Neural Inform Process Syst.* (2018) 31:9492–503. doi: 10.48550/arXiv.1803.01422
41. Huber M. An introduction to causal discovery. *Swiss J Econ Stat.* (2024) 160:14. doi: 10.1186/s41937-024-00131-4
42. Hernán MA, Robins JM. Estimating causal effects from epidemiological data. *J Epidemiol Community Health.* (2006) 60:578–86. doi: 10.1136/jech.2004.029496
43. Hünemann P, Louw B. On the nuisance of control variables in causal regression analysis. *Organ Res Methods.* (2025) 28:138–51. doi: 10.1177/10944281231219274
44. Holland PW. Statistics and causal inference. *J Am Stat Assoc.* (1986) 81:945–60. doi: 10.1080/01621459.1986.10478354
45. Chernozhukov V, Chetverikov D, Demirer M, Duflo E, Hansen C, Newey W, et al. Double/debiased machine learning for treatment and structural parameters. *Econom J* (2018) 21:C1–68. doi: 10.1111/ectj.12097
46. Jung Y, Tian J, Bareinboim E. Estimating identifiable causal effects through double machine learning. In: *Proceedings of the AAAI Conference on Artificial Intelligence.* (2021). Vol. 35. p. 12113–22.
47. Kang JD, Schafer JL. Demystifying double robustness: a comparison of alternative strategies for estimating a population mean from incomplete data (2007).
48. Künzel SR, Sekhon JS, Bickel PJ, Yu B. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proc Natl Acad Sci.* (2019) 116:4156–65. doi: 10.1073/pnas.1804597116
49. Scutari M, Denis J-B. *Bayesian Networks: With Examples in R.* Boca Raton, FL: Chapman and Hall/CRC (2021).
50. Maathuis MH, Kalisch M, Bühlmann P. Estimating high-dimensional intervention effects from observational data (2009).
51. Bareinboim E, Pearl J. Causal inference and the data-fusion problem. *Proc Natl Acad Sci.* (2016) 113:7345–52. doi: 10.1073/pnas.1510507113
52. Rosenbaum PR. Sensitivity analysis for certain permutation inferences in matched observational studies. *Biometrika.* (1987) 74:13–26. doi: 10.1093/biomet/74.1.13
53. VanderWeele TJ, Ding P. Sensitivity analysis in observational research: introducing the e-value. *Ann Intern Med.* (2017) 167:268–74. doi: 10.7326/M16-2607
54. Balke A, Pearl J. Bounds on treatment effects from studies with imperfect compliance. *J Am Stat Assoc.* (1997) 92:1171–6. doi: 10.1080/01621459.1997.10474074
55. Jonzon G, Gabriel EE, Sjölander A, Sachs MC. Adding covariates to bounds: what is the question? (2025).
56. Daza EJ. Causal analysis of self-tracked time series data using a counterfactual framework for n-of-1 trials. *Methods Inf Med.* (2018) 57:e10–e21. doi: 10.3414/ME16-02-0044
57. Zucker DR, Ruthazer R, Schmid CH. Individual (n-of-1) trials can be combined to give population comparative treatment effect estimates: methodologic considerations. *J Clin Epidemiol.* (2010) 63:1312–23. doi: 10.1016/j.jclinepi.2010.04.020
58. Granger CWJ. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica.* (1969) 37(3):424–38. doi: 10.2307/1912791
59. Kim J-M, Lee N, Hwang SY. A copula nonlinear granger causality. *Econ Model.* (2020) 88:420–30. doi: 10.1016/j.econmod.2019.09.052
60. Malinsky D, Spirtes P. Causal structure learning from multivariate time series in settings with unmeasured confounding. In: *Proceedings of 2018 ACM SIGKDD Workshop on Causal Discovery.* PMLR (2018). p. 23–47.
61. Malinsky D, Spirtes P. Learning the structure of a nonstationary vector autoregression. In: *The 22nd International Conference on Artificial Intelligence and Statistics.* PMLR (2019). p. 2986–94.
62. Runge J. Discovering contemporaneous and lagged causal relations in autocorrelated nonlinear time series datasets. In: *Conference on Uncertainty in Artificial Intelligence.* PMLR (2020). p. 1388–97.
63. Sutton RS, Barto AG. *Reinforcement Learning: An Introduction.* Cambridge, MA: MIT Press (2018).
64. Ramos GDO, Bazzan AL. On estimating action regret and learning from it in route choice. In: *Proceedings of the Ninth Workshop on Agents in Traffic and Transportation (ATT-2016).* New York: CEUR-WS.org (2016). p. 1–8.
65. Slivkins A. Introduction to multi-armed bandits. *Found Trends Mach Learn.* (2019) 12:1–286. doi: 10.1561/22000000068
66. Lattimore T, Szepesvári C. *Bandit Algorithms.* Cambridge: Cambridge University Press (2020).
67. Chu W, Li L, Reyzin L, Schapire R. Contextual bandits with linear payoff functions. In: *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics.* JMLR Workshop and Conference Proceedings (2011). p. 208–14.
68. Bareinboim E, Zhang J, Lee S. An introduction to causal reinforcement learning. (2024).
69. Lee S, Bareinboim E. Structural causal bandits: where to intervene? *Adv Neural Inf Process Syst.* (2018) 31:2568–78.
70. Lee S, Bareinboim E. Structural causal bandits with non-manipulable variables. In: *Proceedings of the AAAI Conference on Artificial Intelligence.* (2019). Vol. 33. p. 4164–72.
71. Ghatk G. A change-detection-based thompson sampling framework for non-stationary bandits. *IEEE Trans Comput.* (2020) 70:1670–6. doi: 10.1109/TC.2020.3022634
72. Russo DJ, Van Roy B, Kazerouni A, Osband I, Wen Z. A tutorial on thompson sampling. *Found Trends Mach Learn.* (2018) 11:1–96. doi: 10.1561/22000000070

73. Thompson WR. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*. (1933) 25:285–94. doi: 10.1093/biomet/25.3-4.285
74. Chapelle O, Li L. An empirical evaluation of thompson sampling. *Adv Neural Inf Process Syst*. (2011) 24:2249–57.
75. Fink D. A compendium of conjugate priors (1997). Available at: <http://www.people.cornell.edu/pages/df36/CONJINTRnew%20TEX.pdf>. p. 46.
76. Agrawal S, Goyal N. Analysis of thompson sampling for the multi-armed bandit problem. In: *Conference on Learning Theory*. JMLR Workshop and Conference Proceedings (2012). p. 39–1.
77. Liu X, Deliu N, Chakraborty T, Bell L, Chakraborty B. Thompson sampling for zero-inflated count outcomes with an application to the drink less mobile health study. *arXiv [Preprint]*. *arXiv:2311.14359* (2023).
78. Trovo F, Paladino S, Restelli M, Gatti N. Sliding-window thompson sampling for non-stationary settings. *J Artif Intell Res*. (2020) 68:311–64. doi: 10.1613/jair.1.11407
79. Kalkanli C, Ozgur A. Batched thompson sampling. *Adv Neural Inf Process Syst*. (2021) 34:29984–94.
80. Agrawal S, Goyal N. Thompson sampling for contextual bandits with linear payoffs. In: *International Conference on Machine Learning*. PMLR (2013). p. 127–35.
81. Nelson E, Bhattacharjya D, Gao T, Liu M, Bouneffouf D, Poupard P. Linearizing contextual bandits with latent state dynamics. In: Cussens J, Zhang K, editors. *Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence*. PMLR (2022). Proceedings of Machine Learning Research; vol. 180. p. 1477–87.
82. Bareinboim E, Forney A, Pearl J. Bandits with unobserved confounders: a causal approach. *Adv Neural Inf Process Syst*. (2015) 28:1342–50.
83. Dimakopoulou M, Ren Z, Zhou Z. Online multi-armed bandits with adaptive inference. *Adv Neural Inf Process Syst*. (2021) 34:1939–51.
84. Reimer U, Maier E, Ulmer T. A self-learning application framework for behavioral change support. In: Röcker C, O'Donoghue J, Ziefle M, Helfert M, Molloy W, editors. *Information and Communication Technologies for Ageing Well and e-Health: Second International Conference, ICT4AWE 2016, Rome, Italy, April 21–22, 2016, Revised Selected Papers 2*. Cham: Springer (2017). p. 119–39. doi: 10.1007/978-3-319-62704-5_8
85. MacLean KE. Designing with haptic feedback. In: *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No. 00ch37065)*. IEEE (2000). Vol. 1. p. 783–8.
86. Costa J, Guimbretière F, Jung MF, Choudhury T. Boostmeup: Improving cognitive performance in the moment by unobtrusively regulating emotions with a smartwatch. *Proc ACM Interact Mobile Wearable Ubiquitous Technol*. (2019) 3:1–23. doi: 10.1145/3328911
87. Malenica I, Guo Y, Gan K, Konigorski S. Anytime-valid inference in n-of-1 trials. In: *Machine Learning for Health (MLAH)*. PMLR (2023). p. 307–22.
88. Avina C. The use of self-monitoring as a treatment intervention. In: *Evidence-Based Adjunctive Treatments*. Elsevier (2008). p. 207–19.
89. Hammady R, Arnab S. Serious gaming for behaviour change: a systematic review. *Information*. (2022) 13:142. doi: 10.3390/info13030142
90. Dobkin BH, Dorsch A. The promise of mhealth: daily activity monitoring and outcome assessments by wearable sensors. *Neurorehabil Neural Repair*. (2011) 25:788–98. doi: 10.1177/1545968311425908
91. Jennings HM, Morrison J, Akter K, Kuddus A, Ahmed N, Kumer Shaha S, et al. Developing a theory-driven contextually relevant mhealth intervention. *Glob Health Action*. (2019) 12:1550736. doi: 10.1080/16549716.2018.1550736
92. Daniëls NE, Hochstenbach LM, van Zelst C, van Bokhoven MA, Delespaul PA, Beurskens AJ. Factors that influence the use of electronic diaries in health care: scoping review. *JMIR mHealth and uHealth*. (2021) 9:e19536. doi: 10.2196/19536
93. Mehregany M, Saldivar E. Opportunities and obstacles in the adoption of mHealth. In: *mHealth*. HIMSS Publishing (2020). p. 7–20.
94. Prasad S. Designing for scalability and trustworthiness in mHealth systems. In: *Distributed Computing and Internet Technology: 11th International Conference, ICDCIT 2015, Bhubaneswar, India, February 5–8, 2015. Proceedings 11*. Springer (2015). p. 114–33.
95. Aguiar M, Cejudo A, Epelde G, Chaves D, Trujillo M, Artola G, et al. An approach to boost adherence to self-data reporting in mhealth applications for users without specific health conditions. *BMC Med Inform Decis Mak*. (2025) 25:16. doi: 10.1186/s12911-024-02833-4
96. Kettner F, Schmidt K, Gaedke M. Motivation enhancement in mHealth via gamification. In: *Proceedings of the 2019 4th International Conference on Informatics and Assistive Technologies for Healthcare, Medical Support and Wellbeing (Healthinfo 2019), Valencia, Spain*. (2019). p. 24–8.
97. Tran S, Smith L, El-Den S, Carter S. The use of gamification and incentives in mobile health apps to improve medication adherence: scoping review. *JMIR mHealth uHealth*. (2022) 10:e30671. doi: 10.2196/30671
98. Businelle MS, Hébert ET, Shi D, Benson L, Kezbers KM, Tonkin S, et al. Investigating best practices for ecological momentary assessment: nationwide factorial experiment. *J Med Internet Res*. (2024) 26:e50275. doi: 10.2196/50275
99. Nahum-Shani I, Smith SN, Spring BJ, Collins LM, Witkiewitz K, Tewari A, et al. Just-in-time adaptive interventions (JITAIs) in mobile health: key components and design principles for ongoing health behavior support. *Ann Behav Med*. (2018) 52(6):446–62. doi: 10.1007/s12160-016-9830-8
100. Tackney MS, Cook DG, Stahl D, Ismail K, Williamson E, Carpenter J. A framework for handling missing accelerometer outcome data in trials. *Trials*. (2021) 22:379. doi: 10.1186/s13063-021-05284-8
101. Tackney MS, Stahl D, Williamson E, Carpenter J. Missing step count data? Step away from the expectation–maximization algorithm. *J Meas Phys Behav*. (2022) 5:205–14. doi: 10.1123/jmpb.2022-0002
102. Di J, Demanuele C, Kettermann A, Karahanoglu FI, Cappelleri JC, Potter A, et al. Considerations to address missing data when deriving clinical trial endpoints from digital health technologies. *Contemp Clin Trials*. (2022) 113:106661. doi: 10.1016/j.cct.2021.106661
103. Daza EJ, Hudgens MG, Herring AH. Estimating inverse-probability weights for longitudinal data with dropout or truncation: the xtrccipw command. *Stata J*. (2017) 17:253–78. doi: 10.1177/1536867X1701700202
104. Little RJ, D'agostino R, Cohen ML, Dickersin K, Emerson SS, Farrar JT, et al. The prevention and treatment of missing data in clinical trials. *N Engl J Med*. (2012) 367:1355–60. doi: 10.1056/NEJMsri1203730
105. Nahum-Shani I, Wetter DW, Murphy SA. Adapting just-in-time interventions to vulnerability and receptivity: conceptual and methodological considerations. In: *Digital Therapeutics for Mental Health and Addiction*. Elsevier (2023). p. 77–87.
106. Battalio SL, Conroy DE, Dempsey W, Liao P, Menictas M, Murphy S, et al. Sense2stop: a micro-randomized trial using wearable sensors to optimize a just-in-time-adaptive stress management intervention for smoking relapse prevention. *Contemp Clin Trials*. (2021) 109:106534. doi: 10.1016/j.cct.2021.106534
107. 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) 34:1209. doi: 10.1037/hea0000306
108. Wang L, Miller LC. Just-in-the-moment adaptive interventions (JITAI): a meta-analytical review. *Health Commun*. (2020) 35:1531–44. doi: 10.1080/10410236.2019.1652388
109. Lei H, Lu Y, Tewari A, Murphy SA. An actor-critic contextual bandit algorithm for personalized mobile health interventions. *arXiv [Preprint]*. *arXiv:1706.09090* (2017).
110. Tewari A, Murphy SA. From ads to interventions: contextual bandits in mobile health. In: Rehag J, Murphy S, Kumar S, editors. *Mobile Health: Sensors, Analytic Methods, and Applications*. Cham: Springer (2017). p. 495–517. doi: 10.1007/978-3-319-51394-2_25
111. Trella AL, Zhang KW, Nahum-Shani I, Shetty V, Doshi-Velez F, Murphy SA. Reward design for an online reinforcement learning algorithm supporting oral self-care. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. (2023). Vol. 37. p. 15724–30.
112. Calbimonte J-P, Calvaresi D, Dubosson F, Schumacher M. Towards profile and domain modelling in agent-based applications for behavior change. In: *Advances in Practical Applications of Survivable Agents and Multi-Agent Systems: The PAAMS Collection: 17th International Conference, PAAMS 2019, Ávila, Spain, June 26–28, 2019, Proceedings 17*. Springer (2019). p. 16–28.
113. Bishop CM. *Pattern Recognition and Machine Learning*. New York, NY: Springer (2006).
114. Švihrová R, Marzorati D, Dei Rossi A, Gerosa T, Faraci FD. Data-driven causal discovery: insights from a longitudinal study with wearable data. In: *Proceedings of the 23rd International Conference on Artificial Intelligence in Medicine*. Springer (2025).
115. Velikzhanin A, Wang B, Kwiatkowska M. Bayesian network models of causal interventions in healthcare decision making: literature review and software evaluation. *arXiv [Preprint]*. *arXiv:2211.15258* (2022).
116. Gerstenberg T, Goodman ND, Lagnado DA, Tenenbaum JB. A counterfactual simulation model of causal judgments for physical events. *Psychol Rev*. (2021) 128:936. doi: 10.1037/rev0000281
117. Silva JF, Matos S. Modelling patient trajectories using multimodal information. *J Biomed Inform*. (2022) 134:104195. doi: 10.1016/j.jbi.2022.104195
118. Tian J, Ding F, Wang R, Han G, Yan J, Yuan N, et al. Dynamic trajectory of a patient-reported outcome and its associated factors for patients with chronic heart failure: a growth mixture model approach. *Risk Manag Healthc Policy*. (2022) 31:2083–96. doi: 10.2147/RMHP.S384936
119. Al Osman H, Eid M, El Saddik A. U-biofeedback: a multimedia-based reference model for ubiquitous biofeedback systems. *Multimed Tools Appl*. (2014) 72:3143–68. doi: 10.1007/s11042-013-1590-x
120. Nahum-Shani I, Shaw SD, Carpenter SM, Murphy SA, Yoon C. Engagement in digital interventions. *Am Psychol*. (2022) 77:836. doi: 10.1037/amp0000983
121. Cialdini R. *Influence: Science and Practice*. Boston, MA: Allyn and Bacon (1993).
122. Kaptein M, De Ruyter B, Markopoulos P, Aarts E. Adaptive persuasive systems: a study of tailored persuasive text messages to reduce snacking. *ACM Trans Interact Intell Syst*. (2012) 2:1–25. doi: 10.1145/2209310.2209313

123. Kaptein M, Markopoulos P, de Ruyter B, Aarts E. Personalizing persuasive technologies: explicit and implicit personalization using persuasion profiles. *Int J Hum Comput Stud.* (2015) 77:38–51. doi: 10.1016/j.ijhcs.2015.01.004
124. Rabbi M, Klasnja P, Choudhury T, Tewari A, Murphy S. Optimizing mHealth interventions with a bandit. In: Baumeister H, Montag C, editors. *Digital Phenotyping and Mobile Sensing. Studies in Neuroscience, Psychology and Behavioral Economics.* Cham: Springer (2019). p. 277–91. doi: 10.1007/978-3-030-31620-4_18
125. Liao P, Greenewald K, Klasnja P, Murphy S. Personalized heartsteps: a reinforcement learning algorithm for optimizing physical activity. *Proc ACM Interact Mobile Wearable Ubiquitous Technol.* (2020) 4:1–22. doi: 10.1145/3381007
126. Rabbi M, Aung MH, 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.* (2015). p. 707–18.
127. Spruijt-Metz D, Marlin BM, Pavel M, Rivera DE, Hekler E, De La Torre S, et al. Advancing behavioral intervention and theory development for mobile health: the heartsteps ii protocol. *Int J Environ Res Public Health.* (2022) 19:2267. doi: 10.3390/ijerph19042267
128. Yu S, Nourzad N, Semple RJ, Zhao Y, Zhou E, Krishnamachari B. Careforme: contextual multi-armed bandit recommendation framework for mental health. In: *Proceedings of the IEEE/ACM 11th International Conference on Mobile Software Engineering and Systems.* (2024). p. 92–4.
129. Zhou T, Wang Y, Yan L, Tan Y. Spoiled for choice? Personalized recommendation for healthcare decisions: a multiarmed bandit approach. *Inform Syst Res.* (2023) 34:1493–512. doi: 10.1287/isre.2022.1191
130. Zhu F, Guo J, Li R, Huang J. Robust actor-critic contextual bandit for mobile health (mHealth) interventions. In: *Proceedings of the 2018 ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics.* (2018). p. 492–501.
131. Ghosh S, Guo Y, Hung P-Y, Coughlin LN, Bonar E, Nahum-Shani I, et al. MiWaves reinforcement learning algorithm. *CoRR* [Preprint]. (2024).
132. Xu Z, Jajal H, Choi SW, Nahum-Shani I, Shani G, Psihogios AM, et al. Reinforcement learning on aya dyads to enhance medication adherence. *arXiv* [Preprint]. *arXiv:2502.06835* (2025).
133. Akacha M, Bretz F, Ruberg S. Estimands in clinical trials—broadening the perspective. *Stat Med.* (2017) 36:5–19. doi: 10.1002/sim.7033
134. Huang X, Lin J, Demner-Fushman D. Evaluation of pico as a knowledge representation for clinical questions. In: *AMIA Annual Symposium Proceedings.* American Medical Informatics Association (2006). Vol. 2006. p. 359.
135. Hopewell S, Chan A-W, Collins GS, Hróbjartsson A, Moher D, Schulz KF, et al. Consort 2025 statement: updated guideline for reporting randomised trials. *Lancet.* (2025) 405:1633–40. doi: 10.1016/S0140-6736(25)00672-5
136. Kay R. Statistical principles for clinical trials. *J Int Med Res.* (1998) 26:57–65. doi: 10.1177/030006059802600201
137. Snapinn SM. Noninferiority trials. *Trials.* (2000) 1:1–3. doi: 10.1186/cvm-1-1-019
138. Cuzick J, Sasieni P. Interpreting the results of noninferiority trials—a review. *Br J Cancer.* (2022) 127:1755–9. doi: 10.1038/s41416-022-01937-w
139. Lancaster GA, Thabane L. Guidelines for reporting non-randomised pilot and feasibility studies. *Pilot Feasibility Stud.* (2019) 5:1–6. doi: 10.1186/s40814-019-0499-1
140. Kumar S, Nilsen WJ, Abernethy A, Atienza A, Patrick K, Pavel M, et al. Mobile health technology evaluation: the mHealth evidence workshop. *Am J Prev Med.* (2013) 45:228–36. doi: 10.1016/j.amepre.2013.03.017
141. Abadie A. Using synthetic controls: feasibility, data requirements, and methodological aspects. *J Econ Lit.* (2021) 59:391–425. doi: 10.1257/jel.20191450
142. Collins LM, Kugler KC. *Optimization of Behavioral, Biobehavioral, and Biomedical Interventions.* Cham: Springer International Publishing (2018). Vol. 10. p. 978–3.
143. Dziak JJ, Nahum-Shani I, Collins LM. Multilevel factorial experiments for developing behavioral interventions: power, sample size, and resource considerations. *Psychol Methods.* (2012) 17:153. doi: 10.1037/a0026972
144. Klasnja P, Hekler EB, Shiffman S, Boruvka A, Almirall D, Tewari A, et al. Microrandomized trials: an experimental design for developing just-in-time adaptive interventions. *Health Psychol.* (2015) 34:1220. doi: 10.1037/hea0000305
145. Qian T, Yoo H, Klasnja P, Almirall D, Murphy SA. Estimating time-varying causal excursion effects in mobile health with binary outcomes. *Biometrika.* (2021) 108:507–27. doi: 10.1093/biomet/asaa070
146. Liu X, Qian T, Bell L, Chakraborty B. Incorporating nonparametric methods for estimating causal excursion effects in mobile health with zero-inflated count outcomes. *arXiv* [Preprint]. *arXiv:2310.18905* (2023).
147. Boruvka A, Almirall D, Witkiewitz K, Murphy SA. Assessing time-varying causal effect moderation in mobile health. *J Am Stat Assoc.* (2018) 113:1112–21. doi: 10.1080/01621459.2017.1305274
148. Bell L, Garnett C, Qian T, Perski O, Potts HW, Williamson E, et al. Notifications to improve engagement with an alcohol reduction app: protocol for a micro-randomized trial. *JMIR Res Protoc.* (2020) 9:e18690. doi: 10.2196/18690
149. Golbus JR, Dempsey W, Jackson EA, Nallamotheu BK, Klasnja P. Microrandomized trial design for evaluating just-in-time adaptive interventions through mobile health technologies for cardiovascular disease. *Circ Cardiovasc Qual Outcomes.* (2021) 14:e006760. doi: 10.1161/CIRCOUTCOMES.120.006760
150. Golbus JR, Gupta K, Stevens R, Jeganathan VSE, Luff E, Shi J, et al. A randomized trial of a mobile health intervention to augment cardiac rehabilitation. *NPJ Digit Med.* (2023) 6:173. doi: 10.1038/s41746-023-00921-9
151. Liao P, Klasnja P, Tewari A, Murphy SA. Sample size calculations for micro-randomized trials in mHealth. *Stat Med.* (2016) 35:1944–71. doi: 10.1002/sim.6847
152. Atkinson AC, Biswas A. *Randomised Response-Adaptive Designs in Clinical Trials.* Boca Raton, FL: Chapman & Hall (2019).
153. Hu F, Rosenberger WF. *The Theory of Response-Adaptive Randomization in Clinical Trials.* Hoboken, NJ: John Wiley & Sons (2006).
154. Villar SS, Wason J, Bowden J. Response-adaptive randomization for multi-arm clinical trials using the forward looking gittins index rule. *Biometrics.* (2015) 71:969–78. doi: 10.1111/biom.12337
155. Namkoong H, Daulton S, Bakshy E. Distilled thompson sampling: practical and efficient thompson sampling via imitation learning. *arXiv* [Preprint]. *arXiv:2011.14266* (2020).
156. Almirall D, Nahum-Shani I, Sherwood NE, Murphy SA. Introduction to smart designs for the development of adaptive interventions: with application to weight loss research. *Transl Behav Med.* (2014) 4:260–74. doi: 10.1007/s13142-014-0265-0
157. Nahum-Shani I, Dziak JJ, Wetter DW. Mcmtc: a pragmatic framework for selecting an experimental design to inform the development of digital interventions. *Front Digit Health.* (2022) 4:798025. doi: 10.3389/fdgth.2022.798025
158. McCoy CE. Understanding the intention-to-treat principle in randomized controlled trials. *West J Emerg Med.* (2017) 18:1075. doi: 10.5811/westjem.2017.8.35985
159. Tripepi G, Chesnaye NC, Dekker FW, Zoccali C, Jager KJ. Intention to treat and per protocol analysis in clinical trials. *Nephrology.* (2020) 25:513–7. doi: 10.1111/nep.13709
160. Angrist JD, Gao C, Hull P, Yeh RW. Instrumental variables in randomized trials. *NEJM Evid.* (2025) 4:EVIDctw2400204. doi: 10.1056/EVIDctw2400204
161. Angrist JD, Imbens GW, Rubin DB. Identification of causal effects using instrumental variables. *J Am Stat Assoc.* (1996) 91:444–55. doi: 10.1080/01621459.1996.10476902
162. Peugh JL, Strotman D, McGrady M, Rausch J, Kashikar-Zuck S. Beyond intent to treat (ITT): a complier average causal effect (CACE) estimation primer. *J Sch Psychol.* (2017) 60:7–24. doi: 10.1016/j.jsp.2015.12.006
163. Sussman JB, Hayward RA. An IV for the RCT: using instrumental variables to adjust for treatment contamination in randomised controlled trials. *Bmj.* (2010) 340:c2073. doi: 10.1136/bmj.c2073
164. Berridge DM, Crouchley R. *Multivariate Generalized Linear Mixed Models Using R.* Boca Raton, FL: CRC Press (2011).
165. Goldstein H. *Multilevel Models in Education and Social Research.* New York, NY: Oxford University Press (1987).
166. Dempster AP, Laird NM, Rubin DB. Maximum likelihood from incomplete data via the em algorithm. *J R Stat Soc Ser B Methodol.* (1977) 39:1–22. doi: 10.1111/j.2517-6161.1977.tb01600.x
167. Fong Y, Rue H, Wakefield J. Bayesian inference for generalized linear mixed models. *Biostatistics.* (2010) 11:397–412. doi: 10.1093/biostatistics/kxp053
168. Qian T, Klasnja P, Murphy SA. Linear mixed models with endogenous covariates: modeling sequential treatment effects with application to a mobile health study. *Stat Sci.* (2020) 35:375. doi: 10.1214/19-sts720
169. Chen D-G, Peace K. *Clinical trial data analysis using R* (2010). doi: 10.13140/2.1.3362.1444
170. Qian T, Walton AE, Collins LM, Klasnja P, Lanza ST, Nahum-Shani I, et al. The microrandomized trial for developing digital interventions: experimental design and data analysis considerations. *Psychol Methods.* (2022) 27:874. doi: 10.1037/met0000283
171. Feuerriegel S, Frauen D, Melnychuk V, Schweisthal J, Hess K, Curth A, et al. Causal machine learning for predicting treatment outcomes. *Nat Med.* (2024) 30:958–68. doi: 10.1038/s41591-024-02902-1
172. Häckl S, Koch A, Lasch F. Type-i-error rate inflation in mixed models for repeated measures caused by ambiguous or incomplete model specifications. *Pharm Stat.* (2023) 22:1046–61. doi: 10.1002/pst.2328
173. Amagai S, Pila S, Kaat AJ, Nowinski CJ, Gershon RC. Challenges in participant engagement and retention using mobile health apps: literature review. *J Med Internet Res.* (2022) 24:e35120. doi: 10.2196/35120
174. O'Brien HL, Roll I, Kampen A, Davoudi N. Rethinking (dis) engagement in human-computer interaction. *Comput Hum Behav.* (2022) 128:107109. doi: 10.1016/j.chb.2021.107109

175. Park J, Lee U. Understanding disengagement in just-in-time mobile health interventions. *Proc ACM Interact Mobile Wearable Ubiquitous Technol.* (2023) 7:1–27. doi: 10.1145/3596240
176. Shin G, Feng Y, Jarrahi MH, Gafinowitz N. Beyond novelty effect: a mixed-methods exploration into the motivation for long-term activity tracker use. *JAMIA Open.* (2019) 2:62–72. doi: 10.1093/jamiaopen/ooy048
177. Nahum-Shani I, Rabbi M, Yap J, Philyaw-Kotov ML, Klasnja P, Bonar EE, et al. Translating strategies for promoting engagement in mobile health: a proof-of-concept microrandomized trial. *Health Psychol.* (2021) 40:974. doi: 10.1037/hea0001101
178. Grote T, Berens P. On the ethics of algorithmic decision-making in healthcare. *J Med Ethics.* (2020) 46:205–11. doi: 10.1136/medethics-2019-105586
179. Boudierhem R. Shaping the future of ai in healthcare through ethics and governance. *Humanit Soc Sci Commun.* (2024) 11:1–12. doi: 10.1057/s41599-024-02894-w
180. US Food and Drug Administration. Digital health technologies for remote data acquisition in clinical investigations. *Draft Guidance* (2022).
181. Weissman GE, Mankowitz T, Kanter GP. Unregulated large language models produce medical device-like output. *npj Digit Med.* (2025) 8:148. doi: 10.1038/s41746-025-01544-y
182. Sedgwick P, Greenwood N. Understanding the hawthorne effect. *Bmj.* (2015) 351:h4672. doi: 10.1136/bmj.h4672
183. Rohrer JM. Thinking clearly about correlations and causation: graphical causal models for observational data. *Adv Methods Pract Psychol Sci.* (2018) 1:27–42. doi: 10.1177/2515245917745629
184. Wysocki AC, Lawson KM, Rhemtulla M. Statistical control requires causal justification. *Adv Methods Pract Psychol Sci.* (2022) 5:25152459221095823. doi: 10.1177/25152459221095823
185. Sel K, Hawkins-Daarud A, Chaudhuri A, Osman D, Bahai A, Paydarfar D, et al. Survey and perspective on verification, validation, and uncertainty quantification of digital twins for precision medicine. *npj Digit Med.* (2025) 8:40. doi: 10.1038/s41746-025-01447-y
186. Sulaiman M, Håkansson A, Karlsen R. Fit-twin: a digital twin of a user with wearables and context as input for health promotion (2023).