Investigating the Relationship Between Physical Activity and Tailored Behavior Change Messaging: Connecting Contextual Bandit with Large Language Models

Haochen Song*
fred.song@mail.utoronto.ca
University of Toronto
Toronto, Ontario, Canada

Dominik Hofer*
dominik.Hofer@dhp.lbg.ac.at
Ludwig Boltzmann Institute for Digital Health and Prevention
Salzburg, Austria

Rania Islambouli
Ludwig Boltzmann Institute for Digital Health and Prevention
Salzburg, Austria

Laura Hawkins
University of Toronto
Toronto, Ontario, Canada

Ananya Bhattacharjee
University of Toronto
Toronto, Ontario, Canada

Meredith Franklin
University of Toronto
Toronto, Ontario, Canada

Joseph Jay Williams
University of Toronto
Toronto, Ontario, Canada

ABSTRACT

Machine learning approaches, such as contextual multi-armed bandit (cMAB) algorithms, offer a promising strategy to reduce sedentary behavior by delivering personalized interventions to encourage physical activity. However, cMAB algorithms typically require large participant samples to learn effectively and may overlook key psychological factors that are not explicitly encoded in the model. In this study, we propose a hybrid approach that combines cMAB for selecting intervention types with large language models (LLMs) to personalize message content. We evaluate four intervention types: behavioral self-monitoring, gain-framing, loss-framing, and social comparison, each delivered as a motivational message aimed at increasing motivation for physical activity and daily step count. Message content is further personalized using dynamic contextual factors including daily fluctuations in self-efficacy, social influence, and regulatory focus. Over a seven-day trial, participants receive daily messages assigned by one of four models: cMAB alone, LLM alone, combined cMAB with LLM personalization (cMABxLLM), or equal randomization (RCT). Outcomes include daily step count and message acceptance, assessed via ecological momentary assessments (EMAs). We apply a causal inference framework to evaluate the effects of each model. Our findings offer new insights into the complementary roles of LLM-based personalization and cMAB adaptation in promoting physical activity through personalized behavioral messaging.

*Both authors contributed equally to this research.

KDD 2025 3rd Workshop on Causal Inference and Machine Learning in Practice, August 03–04, 2025, Toronto, ON

ACM ISBN 978-1-4503-XXXX-X/2018/06
https://doi.org/XXXXXXX.XXXXXXX

CCS CONCEPTS

• **Computing methodologies** → *Machine learning approaches*; • **Applied computing** → *Health informatics*; • **Human-centered computing** → *Empirical studies in HCI*.

KEYWORDS

Contextual Multi-Armed Bandit, Large Language Models, Personalized Messaging, Just-In-Time Adaptive Intervention, Physical Activity, Behavior Change, Bayesian Regression, Digital Health

ACM Reference Format:
Haochen Song, Dominik Hofer, Rania Islambouli, Laura Hawkins, Ananya Bhattacharjee, Meredith Franklin, and Joseph Jay Williams. 2025. Investigating the Relationship Between Physical Activity and Tailored Behavior Change Messaging: Connecting Contextual Bandit with Large Language Models. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (KDD 2025 3rd Workshop on Causal Inference and Machine Learning in Practice)*. ACM, New York, NY, USA, 7 pages. https://doi.org/XXXXXXX.XXXXXXX

1 INTRODUCTION

Physical activity, defined as any bodily movement produced by skeletal muscles that requires energy expenditure, plays a crucial role in our health and well-being [29]. Growing evidence links insufficient physical activity to an increased risk of chronic diseases [4, 22, 24] with an estimated four to five million annual deaths due to sedentary lifestyles [16, 23]. In response, a wide array of digital health interventions and activity-promoting applications have emerged [1, 13, 22], designed to support sustained behavior change and encourage regular physical activity across diverse populations.

The effectiveness of such interventions often depend on a participant's shifting daily context and psychological state. To address this, researchers have increasingly adopted the Just-In-Time Adaptive Intervention (JITAI) framework, which delivers supportive messages tailored to users' immediate context and individual needs [11, 25]. Longitudinal studies have shown that such context-aware approaches can help maintain healthy behaviors over time [10, 21].

However, real-time fluctuations in mood, stress, self-efficacy, and social influence continue to moderate intervention effectiveness [5], presenting a key challenge for adaptive personalization.

This paper proposes a hybrid intervention strategy that integrates contextual multi-armed bandits (cMABs) with large language models (LLMs) to address this challenge. cMABs provide a structured and interpretable decision-making framework, enabling dynamic selection of predefined intervention types based on contextual inputs such as user mood or perceived self-efficacy [3]. However, effective adaptation via cMABs often requires large sample sizes, and struggles managing numerous psychological or environmental variables [17, 26].

LLMs offer a complementary capability: flexible linguistic personalization. They can adapt the tone, wording, and framing of intervention messages based on user context, generating more relatable and engaging communication. However, while LLMs are powerful in language generation, their outputs lack consistency, and their internal decision-making is often opaque [6, 18].

We propose a synergistic approach that uses cMABs for transparent, data-driven selection of intervention types and LLMs for nuanced, context-sensitive message personalization. This hybrid approach aims to combine the strengths of both: the adaptability and interpretability of cMABs with the expressive power of LLMs. We hypothesize that this combined method will produce interventions that are psychologically attuned and behaviorally effective in promoting physical activity.

We address two core **Research Questions** (RQs):

- **RQ1**: To what extent do contextual factors and the combined use of cMABs and LLMs influence users' acceptance of intervention messages, compared to standalone or random assignment strategies?
- **RQ2**: How do users' perceived personalization and self-reported evaluations (e.g. relevance, motivational appeal, clarity) predict subsequent motivation to engage in physical activity?

The following sections describe our theoretical foundation, study design, causal modeling approach, expected outcomes, and implications for advancing personalized health interventions.

2 RELATED WORK

Multi-armed bandit (MAB) algorithms have been applied to tailor physical activity (PA) interventions for sedentary adults by responding to real-time contextual cues, such as location, time of day, and recent movement. For instance, the HeartSteps by Klasnja et al. trial demonstrated that context-sensitive prompts produced immediate increases in step counts; however, these gains diminished over time [15]. Other research has reported more sustained benefits — the MyBehavior app, led to users walking an extra ten minutes each day [19]. Švihrová et al. discussed the potential of MABs for real-time personalization in digital health, emphasizing their ability to improve engagement and efficacy through adaptive intervention delivery. They highlighted opportunities in combining MABs with causal inference and digital twins for personalized behavior change [30]. Together, these findings underscore the potential of MABs to encourage physical activity but also highlight the challenge of avoiding intervention fatigue and maintaining engagement over the long term.

Parallel work has explored using large language models (LLMs) to generate motivational messages. Jörke et al. employed GPT-4 with prompt chaining to emulate motivational interviewing during a sixteen-participant technology probe; it upheld a supportive tone in 84 percent of its utterances and favored open-ended questions over unsolicited advice [13]. COM-B–primed LLM tailored responses to users' motivational states and re-ranked outputs for enhanced alignment, which experts rated as more empathetic and actionable in simulated dialogues [12]. In a just-in-time adaptive intervention (JITAI) study conducted by Willms and Liu, ChatGPT rapidly produced thirteen theory-based lesson messages for parents to promote child activity, demonstrating the model's capacity for quick, customizable content creation (although expert review was still necessary) [28]. Previous research in this domain has therefore demonstrated LLMs provide benefits such as personalized, supportive communication and fast content generation. Notably, these studies face limitations such as their brief durations, small sample sizes, absence of real-world behavior outcomes, and need for expert oversight.

Recent investigations have begun to integrate contextual multi-armed bandits (cMABs) with LLMs to create more adaptive health-promotion strategies. Karine and Marlin (2025) introduced a hybrid framework in which a Thompson Sampling bandit suggests intervention options, and an LLM filters these suggestions based on user-stated preferences, leading to improved personalization in a physical activity setting [14]. Alamdari et al. (2024) proposed using LLM-generated synthetic data to warm-start CMABs, thereby reducing initial regret and accelerating the learning process [2]. In another study, Brooks et al. (2024) conceptualized LLMs as stochastic treatment generators and developed GAMBITTS—a Thompson Sampling variant that accounts for uncertainty in the content generated by the model—resulting in lower regret in simulated physical activity interventions [7]. Wang et al. (2025) combined a context-aware bandit with an LLM-driven chatbot to deliver tailored motivational messages in the realm of sleep health, demonstrating enhanced engagement and improved health outcomes compared to non-adaptive baselines [27]. Collectively, these efforts suggest that coupling CMABs with LLMs can yield interventions that are more responsive, highly personalized, and overall more effective in promoting behavior change.

3 EXPERIMENT MODELS

This study cross-compares four experimental models. In each model, we ensure that the information collected from participants is the same so that participants are blinded to the assignment of experimental condition. One of these models is a randomized controlled trial (**RCT**), where the intervention is randomly selected for participants regardless of their provided information. The other three comparison models are described below:

3.1 Contextual Multi-Armed Bandits (cMABs)

CMABs dynamically select interventions based on contextual factors at each time step. Let X_t denote the context observed at time

t, which includes individual factors such as self-efficacy, social influence, and regulatory focus. The intervention (or arm) space is denoted by $\mathcal{A}$, where each arm represents a different intervention type, and $|\mathcal{A}| = K$ is the total number of available interventions. At each time step t, the algorithm selects an arm $A_t \in \mathcal{A}$ based on the current context X_t.

Let t denote the index for each decision time step. An outcome (reward) R_t is observed shortly after the intervention is delivered at time t. Let D_t denote the historical dataset available up to time t, with $D_0 = \emptyset$ and, for $t \geq 1$, $D_t = D_{t-1} \cup \{(X_{t-1}, A_{t-1}, R_{t-1})\}$, consisting of tuples of contexts, chosen arms, and observed rewards. The posterior distribution for the expected reward [8] of arm a given context X_t when $A_t = a$ is expressed as:

$$P(\mu_a \mid X_t, D_t) : \mu_a = X_t^\top \beta + D(a)^\top \gamma, \quad \begin{bmatrix} \beta \\ \gamma \end{bmatrix} \mid D_t \sim \mathcal{MVN}(m_t, S_t).$$

Here, β is a shared context coefficient vector, and γ captures the main effects of each arm. The vector $D(a)$ denotes the one-hot encoding of arm a. The posterior parameters, namely the mean m_t and covariance S_t, are updated based on historical data D_t.

Algorithm 1 Contextual Bandit with Linear Payoff

1: Observe current context X_t
2: Sample parameters from posterior: $\begin{bmatrix} \beta^{(t)} \\ \gamma^{(t)} \end{bmatrix} \sim \mathcal{MVN}(m_t, S_t)$
3: Compute sampled expected rewards for each $a \in \mathcal{A}$: $\hat{\mu}_a^{(t)} = X_t^\top \beta^{(t)} + D(a)^\top \gamma^{(t)}$
4: Select arm $A_t = \arg\max_{a \in \mathcal{A}} \hat{\mu}_a^{(t)}$
5: Apply intervention A_t, observe reward R_t
6: Update dataset: $D_{t+1} = D_t \cup \{(X_t, A_t, R_t)\}$
7: Update posterior parameters m_{t+1}, S_{t+1}
8: Proceed to next participant

3.2 Large Language Models (LLMs)

In use of the LLMs, we hold and use the same context X_t and historical dataset D_t as before, but introduce an additional user-generated free-text description, denoted by $\mathcal{L}_t$, as shown in Algorithm 2. This response captures the user's situational context and personal narrative, offering a closer alignment with users' linguistically sophisticated content.

At each time step, the LLM first selects a base intervention $A_t \in \mathcal{A}$ from the **finite** set of possible intervention types. It then generates a customized intervention message $M_t \in \mathcal{M}$, where $\mathcal{M}$ is an **infinite** set of linguistically adaptive messages derived from the base arm space $\mathcal{A}$. This process is modeled as:

$$M_t \sim f_{\text{LLM}}(X_t, \mathcal{L}_t, D_t \mid A_t)$$

where f_{LLM} is treated as an unobserved (black-box) mapping that produces a tailored message based on the structured context X_t, user's free-text input $\mathcal{L}_t$, historical data D_t, and the selected base intervention A_t. Unlike selecting a fixed message from $\mathcal{A}$, the LLM adapts the intervention's language and framing to the user's current psychological and situational context, though the specific decision-making process is not directly observable.

Algorithm 2 LLM-based Personalized Messaging

1: Observe context X_t and user-written input $\mathcal{L}_t$
2: LLM select base intervention $A_t \in \mathcal{A}$ given X_t.
3: Generate personalized message: $M_t = f_{\text{LLM}}(X_t, \mathcal{L}_t, D_t \mid A_t)$
4: Deliver M_t, observe reward R_t
5: Update dataset: $D_{t+1} = D_t \cup \{(X_t, A_t, R_t)\}$
6: Proceed to next time step

3.3 A Combination of cMABs and LLMs (cMABxLLM)

In our combined approach, the cMABs model first selects the optimal arm A_t at each time step based on the user's context X_t and past data D_t:

$$A_t = \arg\max_{a \in \mathcal{A}} \mu_{a,t}$$

where $\mu_{a,t}$ is the expected reward for arm a at time t given X_t and D_t, as defined previously. The posterior parameters m_t and S_t are also updated at each step as described in Section 3.1.

After the intervention arm A_t is chosen, the LLM generates a personalized message using the user's context, their written response $\mathcal{L}_t$, and observed history:

$$M_t = f_{\text{LLM}}(X_t, \mathcal{L}_t, D_t \mid A_t)$$

The above method is outlined in Algorithm 3.

Algorithm 3 cMABxLLM Combined Assignment

1: Observe current context X_t and user-written input $\mathcal{L}_t$
2: Use cMAB model to select arm: $A_t = \arg\max_{a \in \mathcal{A}} \mu_{a,t}$
3: Generate message: $M_t = f_{\text{LLM}}(X_t, \mathcal{L}_t, D_t \mid A_t)$
4: Deliver M_t to user and observe reward R_t
5: Update dataset: $D_{t+1} = D_t \cup \{(X_t, A_t, R_t)\}$
6: Update posterior parameters m_{t+1}, S_{t+1}
7: Proceed to next time step

4 METHODOLOGY

4.1 Study Design

The research protocol was approved by the University of Toronto Research Ethics Board. The study design follows a 7-day protocol, with days 1 and 7 reserved for pre- and post-study evaluation and interviews, and the intervening 5 days dedicated to daily intervention delivery and assessment. On day 1, participants complete a pre-study questionnaire that includes the Behavioral Regulation in Exercise Questionnaire (BREQ-3) to assess motivation based on Self-Determination Theory (SDT) [9]; Perceived Trust in Automation using the PAICE scale [20]; and socio-demographic information including gender, education level, and digital skill proficiency. On day 7, participants complete a post-study survey focused on their experience with the intervention messages including perceptions of personalization, appropriateness, and comparative evaluation across intervention types.

During the 5-day intervention phase (days 2-6), participants ($N = 5$) complete a brief daily survey and are assigned one of four experimental conditions using micro-randomization. That is, each day, participants are randomly assigned to one of the following models: a **randomized control model (RCT)**, **LLM-only**, **cMAB-only**, or the **combined cMABxLLM model**.

This micro-randomization is repeated independently each day of the study, enabling within-subject comparison across conditions.

Before receiving their assigned motivational message, participants complete a set of ecological momentary assessments (EMAs) capturing dynamic psychological and contextual variables:

- **Mood:** Self-reported current mood.
- **Stress Level:** Perceived ability to manage upcoming tasks.
- **Self-Efficacy:** Confidence in engaging in planned physical activity.
- **Social Influence:** Likelihood of joining others in physical activity if prompted.
- **Regulatory Focus:** Orientation toward growth/achievement (promotion), or safety/avoidance (prevention).
- **Written Narrative:** A brief description of daily events and their impact on mood, stress, or physical activity.

Although mood and stress are collected daily, they are not used in the real-time cMAB or LLM assignment procedures. Instead, they are included in post-hoc analysis to assess any potential confounding or moderating effects of high-frequency contextual factors (those that vary hourly) versus lower-frequency ones (those that change daily).

Each day, participants receive one motivational message tailored by the assigned experimental model.

Figure 1 provides an overview of the main steps and timeline in our study design. We explain various components of the study design below.

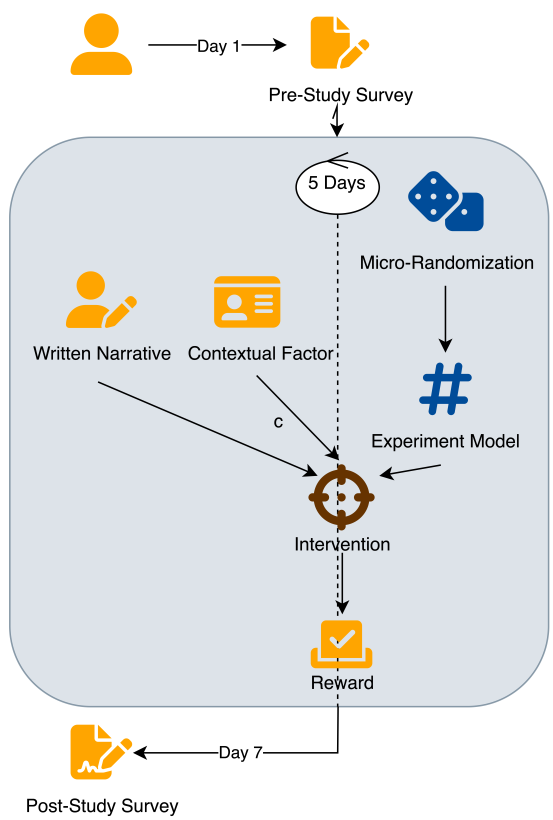

Figure 1: An overview of the study design

4.2 Intervention Messages

Each motivational message corresponds to one of the following four behavior change interventions. When the model does not involve an LLM, a predefined message template corresponding to the assigned type is delivered.

Behavioral Self-Monitoring + Feedback

Great job so far– take a moment to reflect on the time you've spent walking since joining this study. Insert the number of minutes in the box below.

Gain-Framed

Taking a 30-minute walk today could improve your heart health, boost your energy, and elevate your mood for the rest of the evening.

Loss-Framed

Skipping your 30-minute walk today increases your risk of weight gain, poor sleep, and long-term heart health problems.

Social Norms and Comparison Feedback

Many others in your group are meeting their walking goals—join them and keep up the momentum!

4.3 Reward and Feedback Collection

After viewing their assigned intervention message, participants are directed to a response page where they report two key outcomes: 1) Message acceptance – how appropriate, relatable or useful they found the message; and 2) Momentary motivation – their current level of motivation to engage in physical activity.

These responses are treated as the reward signal, which is transmitted to the cMAB training platform (deployed through Javascript on a Google Cloud server). The reward is used to update the cMAB policy for future decision-making, but only in conditions where cMAB is active (i.e. cMAB-onlly or cMABxLLM).

After this, a participants complete a brief feedback survey. This optional qualitative survey invites participants to reflect on the personalization of the message they received, its relevance to their context, and to offer any suggestions for improvement. These narrative responses help us explore perceived personalization but are not used to update the models during the study.

To clarify the variables and written narrative responses, Table 1 provides an overview of which contextual variables are used for model training in each experimental group.

Variable	RCT	cMAB	LLM	cMABxLLM
Mood	×	×	×	×
Stress Level	×	×	×	×
Self-Efficacy	×	✓	✓	✓
Social Influence	×	✓	✓	✓
Regulatory Focus	×	✓	✓	✓
Written Narrative (L_t)	×	×	✓	✓

Table 1: Experiment variables used for each experiment model

4.4 Prompt Example Used in the LLM Model

To generate personalized intervention messages, we provided the LLM with both structured and unstructured user data. The prompt was composed of two parts: a system prompt describing the LLM's task as an assistant and available message types, and as a user prompt containing the participant's current psychological context and self-reflection narrative. Full message templates for each intervention type are described in Section of **Intervention Messages**.

System Prompt

You are an intelligent healthcare assistant tasked with generating personalized health intervention messages to help individuals increase their daily step count. There are four types of intervention available, each defined in the study design:

- Behavioral Self-Monitoring + Feedback
- Gain-Framed Messaging
- Loss-Framed Messaging
- Social Norms & Comparison Feedback

For each participant, you will receive contextual information including self-efficacy, regulatory focus, social influence, and a personal reflection. Based on this information, select one appropriate intervention type and personalize the message using the corresponding template provided in the study design.

User Prompt

Self-efficacy: 72/100 (higher values indicate greater confidence in maintaining physical activity)
Social influence: 64/100 (higher values indicate greater responsiveness to encouragement from others)
Regulatory focus: +3 (positive values indicate gain orientation; negative values indicate loss orientation; range: −6 to +6)
Reflection: "I've been stressed but walking helps clear my mind."

In the LLM-only condition, the model receives all four intervention types and is responsible for selecting and customizing the message. In contrast, in the combined cMABxLLM condition, the intervention type is pre-selected by the cMAB and passed to the LLM, which then generates a message tailored to the participant's psychological context and reflection.

This experimental design allows us to isolate and evaluate the added value of LLM-based message personalization, both independently and in combination with cMAB-driven intervention selection.

5 CAUSAL EFFECT MODELING

To address the **Research Questions** introduced earlier, we construct two distinct causal models—one for each research question. Figure 2 illustrates the hypothesized causal structure of our study. In this diagram, each directed edge represents a potential direct causal relationship.

The experimental condition (G) determines which type of intervention message a participant receives: either a fixed template (Y_F), as used in the RCT and cMAB models, or a personalized message (Y_L)generated by the LLM or cMABxLLM models. Contextual features (X), such as self-efficacy or social influence, inform the generation of Y_L messages but do not influence (Y_F), which are selected from static templates.

Both message types, along with the experimental model G and contextual features X, are expected to influence how participants respond to the message, operationalized as message acceptance (R_1). In turn, message acceptance, together with upstream factors, may influence participants' overall motivation to engage in physical activity (R_2). We also include a set of collected variables (C) such as mood and stress as potential confounders that may affect X or outcomes indirectly. These potential confounders are included in the study to enable post-hoc adjustment, if needed.

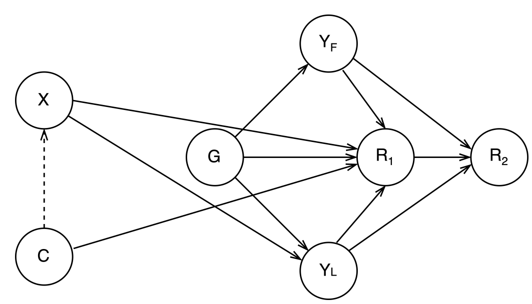

Figure 2: Causal DAG describing the modeling structure of our study.

Variable notations:

- X: Contextual features (e.g., self-efficacy, social influence)
- C: Potential confounders (e.g., mood, stress level)
- G: Experimental model
- Y_F: Fixed message content (RCT and cMAB models)
- Y_L: Personalized message content (LLM and cMABxLLM models)
- R_1: Message acceptance rating

- R_2: Change in motivation toward physical activity (pre- to post-study)

RQ1 evaluates how different intervention models (G) and contextual features (X) influence message acceptance (R_1). To answer this, we compare R_1 outcomes across all four models using offline post-study analysis. The micro-randomization design, in which each participant is randomly assigned to one model per day, mitigates individual-level confounding by distributing conditions evenly within subjects.

We fit a linear mixed-effects model with fixed effects for the intervention model (G) and contextual features (X), and random intercepts for participants to account for repeated measures. Message acceptance (R_1) is treated as a continuous variable on a 5-point Likert scale from 1 (not acceptable) to 5 (liked very much). We hypothesize that fixed-content models (RCT and cMAB-only) will have lower average acceptance than personalized models (LLM and cMABxLLM), although the relative performance of cMABxLLM vs LLM-only is not a priori certain.

To further investigate the role of contextual personalization, particularly in the LLM-personalized arms, we use a Bayesian linear regression model restricted to data from cMAB and cMABxLLM . This model accounts for real-time adaptation and explores how contextual features moderate the efficacy of each intervention type.

$$\text{acceptance\_rating}_i = \alpha + A_i^\top \boldsymbol{\beta}_A + X_i^\top \boldsymbol{\beta}_X + (A_i \otimes X_i)^\top \boldsymbol{\beta}_{AX}$$

where A_i is the intervention type, X_i is the participant's contextual state, and $A_i \otimes X_i$ denotes their interaction terms. This model captures the causal pathways $G \rightarrow Y_F$ & $Y_L \rightarrow R_1$ and $C \dashrightarrow X \rightarrow R_1$ in the DAG (Figure 2). Non-informative priors are used to reflect limited prior knowledge and allow the data to guide inference.

RQ2 examines the longer-term (longitudinal) effects of intervention on participants' motivation to engage in physical activity. Motivation change (R_2) is measured as the difference between post-study and pre-study measures using the BREQ-3 survey [9] (i.e., R_2 = post − pre).

To model this relationship, we use participants' aggregated daily message acceptance ratings (R_1), contextual features X, and confounders C as predictors in a linear mixed-effects regression with motivation change (R_2) as the outcome. The experimental model (G) and baseline characteristics are included as covariates. This approach is aligned with Behavior Change Technique (BCT) and Self-Determination Theory (SDT) and corresponds to the causal pathway ($C \dashrightarrow X \rightarrow R_1 \rightarrow R_2$) in the DAG.

This longitudinal analysis enables us to test whether daily perceptions of personalization and message quality (captured in R_1) contribute meaningfully to longer-term motivation for behavior change, thereby offering insight into how LLM-driven personalization can affect health-related outcomes.

6 PRELIMINARY RESULTS

At this stage, our work remains a small-scale pilot deployment. We have enrolled five participants and started data collection on June 5. The primary goal of this phase is to establish proof of concept and evaluate the combined cMAB×LLM approach against three benchmarks: LLM alone, cMAB alone, and RCT. Although we do not anticipate gathering enough data for extensive training by the initial deadline, we will compile and present an analytical summary as soon as the first batch of data is available. All system-generated contextual information will be fully disclosed. Time permitting, we will also perform sentiment analysis on participants' written feedback and monitor motivational shifts over time. Insights from this pilot will guide the design of a larger study scheduled for July.

7 DISCUSSION

7.1 Contribution

We introduced a new framework that combines contextual multi-armed bandits (cMABs) and large language models (LLMs) for more personalized adaptive intervention assignment. Our main contribution provides a clear, step-by-step guideline for researchers to analyze, implement, and extend this combined cMABxLLM approach. At this stage, the project is in the pilot phase, and there are no empirical results yet. We will share our code and system setup, which are implemented in Google Colab and JavaScript, so others can try, check, and adapt the process for their own studies.

As we collect more data, we plan to use Bayesian penalized regression to reduce overfitting. We also expect that building a reliable model will require a larger dataset. To the best of our knowledge, there is no published study that provides a practical, detailed example of a cMABxLLM pipeline. We hope this work will help other researchers start similar projects and improve on our approach in the future. When more data become available, we will add qualitative interviews and sentiment analysis to learn more about user experience and changes in motivation.

7.2 Limitations

Several limitations to the current study design should be noted. First, the small sample size and short duration of the pilot deployment will limit statistical power and the ability to conclude long-term effects. Meanwhile, cMAB models are sensitive to sparse data early in the study, which can lead to suboptimal assignments and potential bias. In addition, cMAB and LLM models are not trained on identical data: cMAB uses only contextual factors, while LLM models incorporate additional participant information. We welcome discussion and suggestions on these design choices.

The use of language models to generate personalized messages introduces further complexity. The effectiveness of these messages can vary with the quality of the language model and the individual interpretations of the participants, and there remain challenges in model interpretability and the risk of unintended bias. Together, these limitations highlight the need for larger, longer-term studies to validate our findings and further refine adaptive and language-based intervention approaches.

REFERENCES

[1] Adrian Aguilera, Marvyn Arévalo Avalos, Jing Xu, Bibhas Chakraborty, Caroline Figueroa, Faviola Garcia, Karina Rosales, Rosa Hernandez-Ramos, Chris Karr, Joseph Williams, et al. 2024. Effectiveness of a Digital Health Intervention Leveraging Reinforcement Learning: Results From the Diabetes and Mental Health Adaptive Notification Tracking and Evaluation (DIAMANTE) Randomized Clinical Trial. *Journal of medical Internet research* 26 (2024), e60834.

[2] Parand A. Alamdari, Yanshuai Cao, and Kevin H. Wilson. 2024. Jump Starting Bandits with LLM-Generated Prior Knowledge. doi:10.48550/arXiv.2406.19317 arXiv:2406.19317 [cs].

[3] Mawulolo K. Ameko, Miranda L. Beltzer, Lihua Cai, Mehdi Boukhechba, Bethany A. Teachman, and Laura E. Barnes. 2020. Offline Contextual Multi-armed Bandits for Mobile Health Interventions: A Case Study on Emotion Regulation. In *Fourteenth ACM Conference on Recommender Systems*. 249–258. doi:10.1145/3383313.3412244 arXiv:2008.09472 [cs].

[4] Elizabeth Anderson and J Larry Durstine. 2019. Physical activity, exercise, and chronic diseases: A brief review. *Sports medicine and health science* 1, 1 (2019), 3–10.

[5] Ananya Bhattacharjee, Joseph Jay Williams, Jonah Meyerhoff, Harsh Kumar, Alex Mariakakis, and Rachel Kornfield. 2023. Investigating the Role of Context in the Delivery of Text Messages for Supporting Psychological Wellbeing. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–19. doi:10.1145/3544548.3580774

[6] Ananya Bhattacharjee, Sarah Yi Xu, Pranav Rao, Yuchen Zeng, Jonah Meyerhoff, Syed Ishtiaque Ahmed, David C. Mohr, Michael Liut, Alex Mariakakis, Rachel Kornfield, and Joseph Jay Williams. 2025. Perfectly to a Tee: Understanding User Perceptions of Personalized LLM-Enhanced Narrative Interventions. doi:10.48550/arXiv.2409.16732 arXiv:2409.16732 [cs].

[7] Marc Brooks, Gabriel Durham, Kihyuk Hong, and Ambuj Tewari. 2025. Generator-Mediated Bandits: Thompson Sampling for GenAI-Powered Adaptive Interventions. doi:10.48550/arXiv.2505.16311 arXiv:2505.16311 [stat].

[8] Wei Chu, Lihong Li, Lev Reyzin, and Robert Schapire. 2011. Contextual Bandits with Linear Payoff Functions. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*. JMLR Workshop and Conference Proceedings, 208–214. https://proceedings.mlr.press/v15/chu11a.html ISSN: 1938-7228.

[9] Luis Cid, Diogo Monteiro, Diogo Teixeira, Pedro Teques, Susana Alves, João Moutão, Marlene Silva, and António Palmeira. 2018. The Behavioral Regulation in Exercise Questionnaire (BREQ-3) Portuguese-Version: Evidence of Reliability, Validity and Invariance Across Gender. *Frontiers in Psychology* 9 (Oct. 2018), 1940. doi:10.3389/fpsyg.2018.01940

[10] Katharina Feil, Sarah Allion, Susanne Weyland, and Darko Jekauc. 2021. A Systematic Review Examining the Relationship Between Habit and Physical Activity Behavior in Longitudinal Studies. *Frontiers in Psychology* 12 (March 2021). doi:10.3389/fpsyg.2021.626750 Publisher: Frontiers.

[11] Wendy Hardeman, Julie Houghton, Kathleen Lane, Andy Jones, and Felix Naughton. 2019. A systematic review of just-in-time adaptive interventions (JITAIs) to promote physical activity. *International Journal of Behavioral Nutrition and Physical Activity* 16, 1 (April 2019), 31. doi:10.1186/s12966-019-0792-7

[12] Narayan Hegde, Madhurima Vardhan, Deepak Nathani, Emily Rosenzweig, Cathy Speed, Alan Karthikesalingam, and Martin Seneviratne. 2024. Infusing behavior science into large language models for activity coaching. *PLOS Digital Health* 3, 4 (April 2024), e0000431. doi:10.1371/journal.pdig.0000431

[13] Matthew Jörke, Shardul Sapkota, Lyndsea Warkenthien, Niklas Vainio, Paul Schmiedmayer, Emma Brunskill, and James A Landay. 2025. GPTCoach: Towards LLM-Based Physical Activity Coaching. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–46.

[14] Karine Karine and Benjamin M. Marlin. 2025. Combining LLM decision and RL action selection to improve RL policy for adaptive interventions. doi:10.48550/arXiv.2501.06980 arXiv:2501.06980 [cs].

[15] Predrag Klasnja, Shawna Smith, Nicholas J. Seewald, Andy Lee, Kelly Hall, Brook Luers, Eric B. Hekler, and Susan A. Murphy. 2019. Efficacy of Contextually Tailored Suggestions for Physical Activity: A Micro-randomized Optimization Trial of HeartSteps. *Annals of Behavioral Medicine: A Publication of the Society of Behavioral Medicine* 53, 6 (May 2019), 573–582. doi:10.1093/abm/kay067

[16] I-Min Lee, Eric J Shiroma, Felipe Lobelo, Pekka Puska, Steven N Blair, and Peter T Katzmarzyk. 2012. Effect of physical inactivity on major non-communicable diseases worldwide: an analysis of burden of disease and life expectancy. *The lancet* 380, 9838 (2012), 219–229.

[17] Huitian Lei, Yangyi Lu, Ambuj Tewari, and Susan A. Murphy. 2022. An Actor-Critic Contextual Bandit Algorithm for Personalized Mobile Health Interventions. doi:10.48550/arXiv.1706.09090 arXiv:1706.09090 [stat].

[18] Dillon Plunkett, Adam Morris, Keerthi Reddy, and Jorge Morales. 2025. Self-Interpretability: LLMs Can Describe Complex Internal Processes that Drive Their Decisions, and Improve with Training. doi:10.48550/arXiv.2505.17120 arXiv:2505.17120 [cs].

[19] Mashfiqui Rabbi, Angela Pfammatter, Mi Zhang, Bonnie Spring, and Tanzeem Choudhury. 2015. Automated Personalized Feedback for Physical Activity and Dietary Behavior Change With Mobile Phones: A Randomized Controlled Trial on Adults. *JMIR mHealth and uHealth* 3, 2 (May 2015), e42. doi:10.2196/mhealth.4160

[20] Teresa Scantamburlo, Atia Cortés, Francesca Foffano, Cristian Barrué, Veronica Distefano, Long Pham, and Alessandro Fabris. 2023. Artificial Intelligence across Europe: A Study on Awareness, Attitude and Trust. doi:10.48550/arXiv.2308.09979 arXiv:2308.09979 [cs].

[21] Ali Shameli, Tim Althoff, Amin Saberi, and Jure Leskovec. 2017. How Gamification Affects Physical Activity: Large-scale Analysis of Walking Challenges in a Mobile Application. doi:10.48550/arXiv.1702.07437 arXiv:1702.07437 [cs].

[22] Ben Singh, Timothy Olds, Rachel Curtis, Dorothea Dumuid, Rosa Virgara, Amanda Watson, Kimberley Szeto, Edward O'Connor, Ty Ferguson, Emily Eglitis, et al. 2023. Effectiveness of physical activity interventions for improving depression, anxiety and distress: an overview of systematic reviews. *British journal of sports medicine* 57, 18 (2023), 1203–1209.

[23] Tessa Strain, Søren Brage, Stephen J Sharp, Justin Richards, Marko Tainio, Ding Ding, Jacques Benichou, and Paul Kelly. 2020. Use of the prevented fraction for the population to determine deaths averted by existing prevalence of physical activity: a descriptive study. *The Lancet Global Health* 8, 7 (2020), e920–e930.

[24] Jane Thornton, Taniya Nagpal, Kristen Reilly, Moira Stewart, and Robert Petrella. 2022. The 'miracle cure': how do primary care physicians prescribe physical activity with the aim of improving clinical outcomes of chronic disease? A scoping review. *BMJ Open Sport & Exercise Medicine* 8, 3 (2022), e001373.

[25] Claire R. van Genugten, Melissa S. Y. Thong, Wouter van Ballegooijen, Annet M. Kleiboer, Donna Spruijt-Metz, Arnout C. Smit, Mirjam A. G. Sprangers, Yannik Terhorst, and Heleen Riper. 2025. Beyond the current state of just-in-time adaptive interventions in mental health: a qualitative systematic review. *Frontiers in Digital Health* 7 (Jan. 2025). doi:10.3389/fdgth.2025.1460167 Publisher: Frontiers.

[26] Yogatheesan Varatharajah and Brent Berry. 2022. A Contextual-Bandit-Based Approach for Informed Decision-Making in Clinical Trials. *Life* 12, 8 (Aug. 2022), 1277. doi:10.3390/life12081277

[27] Xingbo Wang, Janessa Griffith, Daniel A. Adler, Joey Castillo, Tanzeem Choudhury, and Fei Wang. 2025. Exploring Personalized Health Support through Data-Driven, Theory-Guided LLMs: A Case Study in Sleep Health. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–15. doi:10.1145/3706598.3713852 arXiv:2502.13920 [cs].

[28] Amanda Willms and Sam Liu. 2024. Exploring the Feasibility of Using ChatGPT to Create Just-in-Time Adaptive Physical Activity mHealth Intervention Content: Case Study. *JMIR medical education* 10 (Feb. 2024), e51426. doi:10.2196/51426

[29] World Health Organization. 2024. Physical Activity. https://www.who.int/news-room/fact-sheets/detail/physical-activity Accessed: 2025-05-29.

[30] Radoslava Švihrová, Alvise Dei Rossi, Davide Marzorati, Athina Tzovara, and Francesca Dalia Faraci. 2025. Designing digital health interventions with causal inference and multi-armed bandits: a review. *Frontiers in Digital Health* 7 (June 2025), 1435917. doi:10.3389/fdgth.2025.1435917