

Exploring Personalized Health Support through Data-Driven, Theory-Guided LLMs: A Case Study in Sleep Health

Xingbo Wang
Weill Cornell Medicine,
Cornell University
New York, NY, USA
xiw4011@med.cornell.edu
wangxbzb@gmail.com

Janessa Griffith
Cornell Tech
New York, NY, USA
janessa.griffith@cornell.edu

Daniel A. Adler
Cornell Tech
New York, NY, USA
daa243@cornell.edu

Joey Castillo
Cornell Tech
New York, NY, USA
jc3292@cornell.edu

Tanzeem Choudhury
Cornell Tech
New York, NY, USA
tanzeem.choudhury@cornell.edu

Fei Wang
Weill Cornell Medicine
New York, NY, USA
few2001@med.cornell.edu

ABSTRACT

Despite the prevalence of sleep-tracking devices, many individuals struggle to translate data into actionable improvements in sleep health. Current methods often provide data-driven suggestions but may not be feasible and adaptive to real-life constraints and individual contexts. We present HEALTHGURU, a novel large language model-powered chatbot to enhance sleep health through data-driven, theory-guided, and adaptive recommendations with conversational behavior change support. HEALTHGURU's multi-agent framework integrates wearable device data, contextual information, and a contextual multi-armed bandit model to suggest tailored sleep-enhancing activities. The system facilitates natural conversations while incorporating data-driven insights and theoretical behavior change techniques. Our eight-week in-the-wild deployment study with 16 participants compared HEALTHGURU to a baseline chatbot. Results show improved metrics like sleep duration and activity scores, higher quality responses, and increased user motivation for behavior change with HEALTHGURU. We also identify challenges and design considerations for personalization and user engagement in health chatbots.

CCS CONCEPTS

• **Human-centered computing** → **Interactive systems and tools; HCI design and evaluation methods; Natural language interfaces**; • **Computing methodologies** → *Natural language processing*.

KEYWORDS

Personalized Health, Sleep Health, Behavior Change, Large Language Models, Activity Recommendations

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI '25, April 26-May 1, 2025, Yokohama, Japan

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-1394-1/25/04

<https://doi.org/10.1145/3706598.3713852>

ACM Reference Format:

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 *CHI Conference on Human Factors in Computing Systems (CHI '25)*, April 26-May 1, 2025, Yokohama, Japan. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3706598.3713852>

1 INTRODUCTION

Sleep is a fundamental pillar of human health, profoundly influencing physical well-being, cognitive function, and emotional resilience [30]. With modern wearable devices like smartwatches and rings, monitoring sleep patterns has become more accessible than ever. These technologies provide detailed metrics—sleep duration, stages, heart rate variability, and more—offering valuable insights into nightly rest. However, they often fall short in delivering actionable guidance. Users are frequently presented with an abundance of metrics but lack clear guidance on how to interpret and apply this data to improve their sleep quality.

Many systems have aimed to bridge this gap through personalized sleep interventions, but most have shown significant limitations. Commercial devices like Fitbit, Whoop, and the Oura Ring, as well as research prototypes like SleepCoach [10], provide detailed insights into sleep patterns by correlating sleep factors with sleep quality but often offer generic recommendations that rely on predefined static templates. These templates are not flexible enough to adapt to users' evolving preferences, schedules, or environmental factors. SleepGuru [25] incorporates users' calendars to predict sleep pressure and recommends sleep schedules accordingly. However, it falls short in addressing crucial real-life variables that extend beyond calendar events, such as physical activity or changes in personal habits.

Conversational chatbots have emerged as a promising interaction paradigm for eliciting user personal contexts and preferences and delivering personalized health interventions [5, 12]. However, many health-focused chatbots reply on template-based or rule-based methods, limiting their conversational flexibility and their ability to provide dynamic personalization [24]. This rigidity can lead to repetitive or irrelevant suggestions, potentially reducing engagement and motivation for behavior change over time.

Recent advances in large language models (LLMs) have shown great potential for enabling more flexible, context-aware conversations. However, most existing LLM applications in health are still in their early stages. They struggle to natively support processing and interpreting raw and complex time-series sensor data, as well as translating data insights into actionable plans tailored to users' personal contexts. This gap hinders their effectiveness in driving sustained behavior change in a structured and systematic manner.

To address these identified gaps, we present HEALTHGURU, a novel LLM-powered chatbot augmented by a multi-agent framework to provide personalized, data-driven, and theory-guided sleep health support. HEALTHGURU seamlessly integrates wearable device data, contextual information, and established behavior change theories to deliver adaptive recommendations and motivational support. To tackle the challenge of delivering context-aware, adaptive interventions, we first build a contextual multi-armed bandit (MAB) model that dynamically learns and suggests sleep-enhancing activities by considering real-time environment factors (e.g., time and weather) and individual physiological sleep data collected from wearable devices. Informed by behavioral theories [3, 23], this model balances exploiting previously successful sleep-enhancing activities and exploring new, potentially beneficial options. This ensures that recommendations remain relevant and adapt to users' evolving contexts and preferences. To enhance real-time adaptability and intervention effectiveness, we leverage LLMs that incorporate behavior change techniques to deliver the activity recommendations in a natural and motivational conversation. Combining the strengths of LLMs and the contextual MAB model while overcoming their limitations, we design a multi-agent framework where specialized agents coordinate LLMs' conversational and reasoning capabilities with context-aware dynamic recommendations. These agents work collaboratively to interpret and integrate wearable data insights with context information and deliver adaptive recommendations tailored to users' unique circumstances in a theory-guided conversation. To evaluate HEALTHGURU, we conduct an eight-week in-the-wild deployment study with 16 participants, comparing our system to a baseline chatbot without context-aware recommendations or theory-guided conversations. Our study adopts a within-subjects design. The results show significant improvements in sleep metrics (e.g., longer sleep duration) and better activity scores, when using HEALTHGURU compared to the baseline. Additionally, participants report higher levels of motivation to engage in sleep-enhancing behaviors, attributed to the personalized and context-aware recommendations provided by the system.

The major contributions of this work are:

- We introduce a novel LLM-powered chatbot with a multi-agent framework that integrates wearable data analysis, contextual factors, and a multi-armed bandit model to deliver personalized, timely, and adaptive sleep recommendations guided by behavior change theories.
- We conduct an eight-week in-the-wild deployment study to evaluate the effectiveness and usability of HEALTHGURU compared to a baseline, showing improvements through sleep metrics and user feedback.
- We provide design implications for future personalized health system designs and LLM applications in health applications.

2 BACKGROUND AND RELATED WORK

2.1 Sleep Sensing and Tracking

Lullaby [16] uses multiple sensors (e.g., sound, temperature) to assess and visualize environmental and sleep data together, helping users understand their sleep. WAKE [38] is a behind-the-ear device that detects microsleep. FitBit¹ tracks sleep duration and stages (light, deep, REM) and provides overall sleep quality assessments. Whoop² monitors fitness, sleep, and recovery metrics, offering recommendations to improve human performance. The Oura Ring³ collects precise body signals (e.g., temperature, heart rate variability, movement) to generate readiness, sleep, and activity scores, along with personalized suggestions. We chose the Oura Ring for our research due to its high-quality sleep measurements and activity detection, which align well with our goals for personalized sleep-enhancing recommendations.

2.2 Sleep Feedback and Recommendation

Sleep feedback and recommendation systems generally focus on two key aspects: 1) improving sleep hygiene awareness and 2) providing personalized recommendations based on user lifestyles to improve sleep health.

To improve sleep awareness, systems use various visualization approaches. ShutEye [4] displays activity timelines on mobile wallpapers to indicate optimal timing for sleep-impacting activities. SleepExplorer [28] is a web-based tool that uses data visualizations to show the impacts of different contextual factors on sleep quality, such as steps, calories, and water consumed. SleepTight [6] is a mobile application that further supports self-reporting on influential sleep factors and their impacts on sleep.

To provide personalized sleep recommendations, many studies incorporate user modeling techniques that consider user characteristics, behaviors, and contexts. SleepCoach [10] utilizes static user models to correlate sleep factors with predefined recommendation templates of general guidelines, such as "For the next 6 days, try going to bed at 11 PM." Daskalova et al. [9] improved personalization by clustering users into groups based on shared physical profiles or recommendations. However, this cohort-based model is inherently static and may overlook the nuances of individual behaviors and preferences over time. Lee et al. [25] built a dynamic model that considers users' daily physiological data and calendar events to predict sleep pressure and sleepiness, which are used to recommend sleep schedules. However, it only focuses on calendar events, which constrains their applicability to broader lifestyle contexts and users' preferences. Building upon these studies, we adopt a dynamic user modeling approach that incorporates just-in-time adaptive intervention (JITAI) principles [35]. JITAI is a design methodology that aims to deliver timely and personalized support by dynamically adjusting to users' internal and contextual conditions when individuals are in need of health support and receptive to engaging with an intervention. To operationalize JITAI for sleep health, we build a contextual multi-armed bandit model inspired

¹<https://www.fitbit.com/global/us/technology/fitbit-app>

²<https://www.whoop.com/us/en/thelocker/introducing-whoop-coach-powered-by-openai/>

³<https://ouraring.com/oura-experience>

by previous reinforcement learning approaches [29, 41] that personalizes interventions by learning from user behaviors over time. The model, informed by self-efficacy [3] and habit formation [23] theories, dynamically adjusts recommendations by balancing the exploitation of proven successful sleep-enhancing activities with the exploration of novel ones. Our model considers both immediate contextual factors (e.g., time and weather) and long-term user physiological data from wearable devices. Furthermore, to enhance real-time adaptability and intervention effectiveness, we deliver these recommendations through a theory-guided LLM-based chatbot, which integrates behavior change techniques to address users' needs and preferences.

2.3 LLMs for Health Applications

Kim et al. [18] leveraged LLMs to streamline natural language interactions that can understand users' utterances and automatically generate questions for users to co-build retrospective activity logs for self-tracking. Englhardt et al. [11] systematically investigated the use and evaluation of LLMs for multi-sensor data classification and reasoning to facilitate the clinical workflows of therapists. Abbasian et al. [1] presented a personalized LLM-powered framework for conversational health agents. The framework aims to deliver personalized responses to users' healthcare-related queries via the analyses of users' questions and relevant and essential data from different external sources. Recently, PH-LLM [7], a version of Gemini, has been fine-tuned for personal health applications in sleep and fitness, demonstrating performance comparable to human experts on long-form case studies and professional exams. PHIA [31] is an LLM-based agent framework that uses iterative reasoning, code generation, and web search to analyze and interpret personal health data from wearables. However, all these LLMs and systems focus on static screenshots of user data. Building upon prior work, we build an LLM-powered chatbot. This chatbot integrates an activity recommendation model that analyzes dynamic users' sleep and activity patterns and offers personalized advice or answers specific questions about their sleep health. Moreover, by employing a behavior change framework, the chatbot motivates users to adopt recommended activities, thereby fostering better sleep habits.

3 HEALTHGURU

HEALTHGURU is a novel personalized health support chatbot designed to promote behavior change for better sleep health. By integrating wearable data, contextual information, and established behavior change theories, it offers data-driven, theory-guided support through an LLM-based multi-agent conversational interface.

3.1 Data-Driven and Theory-Guided Health Support

In designing HEALTHGURU, we focused on four key requirements:

- **Provide quantitative health data analytics.** Wearable devices facilitate health monitoring by providing rich, longitudinal data streams [39]. HEALTHGURU integrates relevant wearable data to provide a holistic view of users' health status, offering tangible metrics to track sleep and activities while providing nuanced, data-driven health advice.
- **Adapt to personal contexts and preferences.** Health behaviors are deeply influenced by environmental and personal contexts [42]. Static recommendations often fail to capture the dynamic nature of users' lives [45]. Guided by just-in-time adaptive intervention (JITAI) principles, our system delivers personalized and timely support by considering users' realtime internal states (e.g., physiological metrics) and external contexts (e.g., environment) [35]. We employ a contextual multi-armed bandit approach to dynamically adapt recommendations to users' evolving health states (e.g., physiological metrics) and external contexts (e.g., time, weather). This ensures the relevance and effectiveness of health interventions in changing environments.
- **Support theory-informed motivational practice.** While providing health information and knowledge is essential, it is often insufficient on its own to drive sustained behavior change [17]. Established behavioral change theories [2, 32] offer structured approaches to understanding and influencing health behaviors. Therefore, we ground our system-generated advice on the theoretical techniques and frameworks (e.g., positive reinforcement and goal-setting) to motivate users in their health improvement journey.
- **Enable natural conversational interaction.** The mode of interaction significantly impacts user engagement and the effectiveness of digital health interventions [5]. Conversational interactions have been shown to provide human-like interaction, and increase user satisfaction, trust, and adherence in health applications [20, 24]. Through an LLM-based chatbot, HEALTHGURU facilitates nature and engaging interactions that can deliver JITAI to address users' immediate sleep health concerns and improve user receptivity to interventions according to user feedback.

3.2 System and Data

HEALTHGURU is built on a multi-agent architecture (Figure 1) that processes user inputs, integrates contextual data, and generates personalized, theory-guided responses to promote behavior change and improve sleep health. This section outlines the system workflow and the data types used for personalized feedback.

3.2.1 System Workflow. The system processes user messages and chat history by engaging specialized agents that consider behavior change theory, wearable data, context data, and activity recommendations. When a user sends a message (e.g., "What do you recommend me to do?"), the agent coordinator selects the appropriate agents. For example, the behavior change agent identifies motivational techniques, while the recommendation agent uses current context (e.g., sunny weather at 81°F) to suggest suitable activities. The response agent then compiles these inputs into a concise, coherent reply (as shown in Figure 1, where a late afternoon run is suggested to avoid peak heat, with benefits explained).

3.2.2 Data Collection & Preparation. As summarized in Table 1, we collect three primary types of wearable data for sleep health analysis and activity recommendation: physiological, sleep, and activity data. The selection of these data types is based on their direct relevance to both sleep quality and physical activity and their widespread

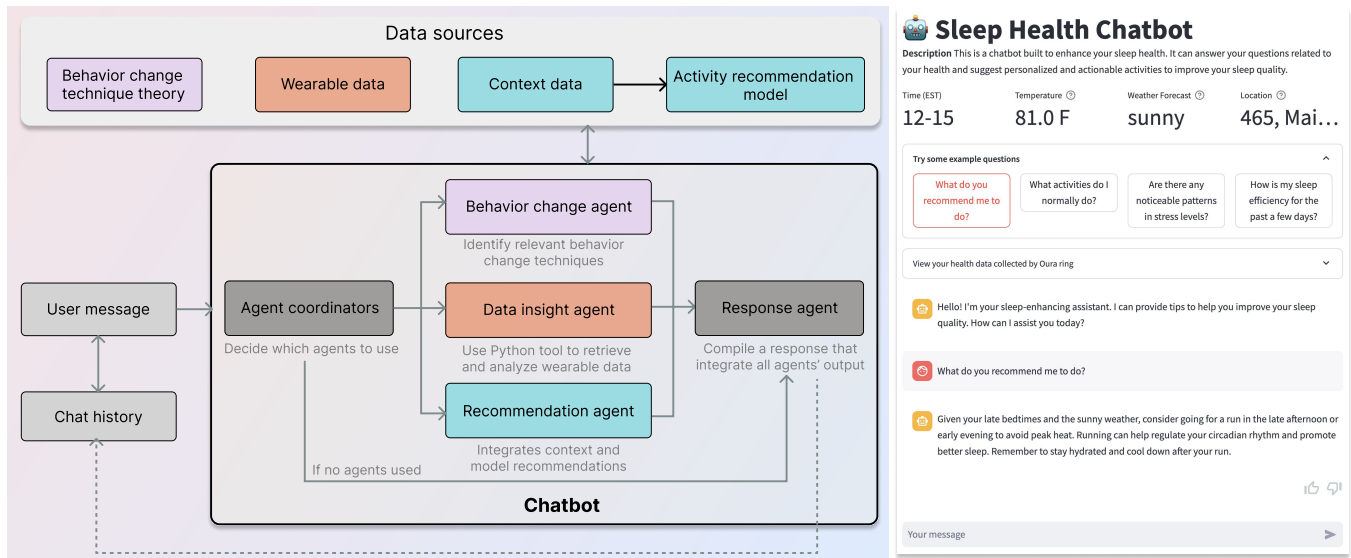


Figure 1: Overview of HEALTHGURU system. The system framework (left) features a multi-agent architecture, incorporating behavior change techniques, wearable data analysis, and context-aware recommendations. The user interface (right) engages users in multi-turn conversations to gain insights into their sleep and activity data and receive personalized advice.

availability in existing commercial wearable devices (e.g., Oura ring, Fitbit, and Whoop). Moreover, the data selection is supported by extensive research in sleep science [22, 36, 37, 44] and human-computer interaction [27]. Physiological data (e.g., HRV) [44] offer quantitative insights into sleep quality and activity intensity. Sleep data details sleep patterns, efficiency, and quality [36]. Activity data captures the type, intensity, and duration of physical activities [22]. We also include context data for generating activity recommendations (in subsection 3.3) and behavior change theory components (in subsection 3.4) for formulating responses.

3.3 Physical Activity Recommendations

We present a novel contextual multi-arm bandit (MAB) model (Figure 2) to recommend personalized activities to improve health outcomes. This approach emphasizes adaptability to user behavior and environmental context factors while maintaining a balance between familiar and novel experiences to motivate behavior change.

Context consideration. We consider the time of day, temperature, and weather [13–15, 19]. These factors are crucial in determining the activity feasibility and effectiveness. For example, the time of day affects user schedules and energy levels. Temperature impacts the comfort of outdoor activities.

Problem formulation. Framed as an optimization task, the system aims to maximize sleep quality scores (measured by devices like the Oura ring) by recommending activities (e.g., Gym, Walking, Yoga, Reading, Meditation) based on current contexts.

Motivation of contextual multi-arm bandit model for activity recommendation is guided by both theoretical and practical considerations. Self-efficacy theory [3] suggests users are more likely to engage in activities they believe they can successfully perform based on their past experience, while habit formation theory [23] emphasizes the importance of consistent positive outcomes.

Practically, the system must address the “cold start” problem of limited initial user data and adapt to changing user preferences and environmental context over time. It also needs to balance between recommending known effective and introducing novel options to prevent user boredom and discover new beneficial behaviors.

To address these requirements, we adopt a contextual multi-arm bandit (MAB) model to guide our activity recommendations. In this framework, each “arm” represents a physical activity and the “reward” is the corresponding sleep quality, while additional context (e.g., time, location, weather) informs decision-making. The algorithm balances exploitation (selecting proven activities) with exploration (trying new or less certain ones), aligning with behavioral theories by recommending feasible, context-appropriate options and fostering habit formation through novel experiences. The model’s adaptive learning capability allows it to respond to changing user preferences and environmental conditions, while also addressing the cold start problem by quickly learning from limited initial data.

Implementation of contextual MAB. We implement the Linear Upper Confidence Bound (LinUCB) algorithm [26] for its ability to handle context information and enable online learning with strong theoretical performance. The core idea behind LinUCB is to model the expected reward of each action as a linear function of the contextual features and incorporate an uncertainty term to balance exploration and exploitation. Considering both estimated rewards and confidence in the estimates, the model selects the actions that have the highest potential for reward. The implementation of LinUCB is described as follows:

- **Context vector construction:** For each recommendation at time t , we create a context vector $x_t \in \mathbb{R}^d$, where d is the dimensionality of the context features. The context vector




Figure 2: Contextual MAB algorithm workflow. The algorithm generates activity recommendations by balancing the exploration and exploitation benefits of activities based on the current context and the corresponding activity history and sleep scores. After the recommendation, the sleep quality for the recommended activities will provide feedback to update the model.

Table 1: Collected wearable data types and definitions.

Category – Measurement	Definition
<i>Physiological Data</i>	
Average Heart Rate	Mean heart rate (in bpm) measured during sleep or activity to assess cardiovascular health.
Lowest Heart Rate	The lowest heart rate (in bpm) recorded during sleep—often occurring in deep sleep stages.
Average HRV	Mean heart rate variability (in ms) during sleep; higher values (e.g., 50–100 ms) suggest better stress resilience.
Stress Level	A computed metric based on physiological markers (such as HRV), typically scaled (e.g., 1–100).
<i>Sleep Data</i>	
Bedtime Start & End	Timestamps indicating when sleep begins and ends (e.g., 2024-07-26T00:26:28-04:00).
Day	The date corresponding to the sleep data (e.g., 2024-07-26).
Sleep Efficiency	Ratio of sleep time to time in bed, expressed as a percentage (1–100).
Readiness Score	A daily readiness score (1–100) derived from sleep quality and physiological measures.
Time in Bed	Total time spent in bed (in seconds), including periods of wakefulness.
Total Sleep Duration	Actual sleep time excluding wake periods (in seconds).
Sleep Score	Overall sleep quality score (1–100) based on duration, efficiency, and physiological data.
<i>Activity Data</i>	
Activity Type	Type of physical activity (e.g., Running, Walking, Cycling).
Intensity	Activity intensity level (e.g., easy, moderate, hard).
Day	The date on which the activity was recorded.
Start & End Time	Timestamps marking the beginning and end of the activity period.

incorporates relevant factors, including time of day, temperature, and weather conditions. Time of day is discretized into seven intervals based on hours: '0-6', '6-9', '9-12', '12-15', '15-18', '18-21', and '21-24'. Temperature values are converted into 'cold', 'mild', 'warm', and 'hot'. Weather is categorized

into 'sunny', 'rain', 'clear', 'windy', and 'snow'. These features are one-hot encoded and concatenated to form the context vector.

- **Action set initialization:** For each user, we define a set \mathcal{A} of possible sleep-enhancing activities (e.g., evening walk, meditation, reading).
- **Decision rule:** At each time step t , for a given context x_t , LinUCB selects the action a_t that maximizes:

$$a_t = \arg \max_{a \in \mathcal{A}} \left(\hat{\theta}_a^\top x_t + \alpha \sqrt{x_t^\top A_a^{-1} x_t} \right),$$

where:

- The term $\hat{\theta}_a^\top x_t$ represents the estimated expected reward for action a given context x_t . $\hat{\theta}_a \in \mathbb{R}^d$ is the estimated parameter vector for action a , representing the learned relationship between the context and the reward (sleep quality) for that action.
- The term $\sqrt{x_t^\top A_a^{-1} x_t}$ represents the uncertainty (standard deviation) in the estimate, forming the upper confidence bound. $A_a \in \mathbb{R}^{d \times d}$ is the covariance matrix (or precision matrix) for action a , initialized as a $d \times d$ identity matrix and updated over time. It records the certainty of our estimates.
- $\alpha (>0)$ is a positive scalar that controls the trade-off between exploration and exploitation. A higher α encourages more exploration.
- **Action execution:** The system recommends the selected activity a_t to the user.
- **Reward observation:** After recommendation, the system observes a reward r_t (e.g., a normalized sleep score between 0 and 1).
- **Update procedure:** We then update the parameters for the suggested action a_t based on the observed reward r_t :
 - (1) Update the covariance matrix by incorporating context information at time t to refine the certainty in parameter estimates: $A_{a_t} = A_{a_t} + x_t x_t^\top$.
 - (2) Update the accumulated reward vector for action a_t : $\mathbf{b}_{a_t} = \mathbf{b}_{a_t} + r_t x_t$.
 - (3) Recompute the parameter vector for action a_t : $\hat{\theta}_{a_t} = A_{a_t}^{-1} \mathbf{b}_{a_t}$.

3.4 Sleep Health Chatbot

While quantitative data from users (e.g., sleep, activity, context) is valuable for promoting sleep health, it often overlooks nuanced personal factors such as lifestyle preferences and individual barriers. To bridge this limitation, we develop an LLM-powered chatbot that enables natural, flexible interactions. By integrating quantitative evidence with behavior change frameworks, the chatbot delivers data-driven and theory-informed advice tailored to users' unique circumstances. This integrated approach aims to create a positive feedback loop where recommended behavior changes lead to improved sleep outcomes.

3.4.1 Build Theoretical Framework for Conversation. We incorporate widely recognized behavior change theories to inform the chatbot's conversational flow and recommendation strategies. The use of these theories provides a structured and solid foundation for identifying the determinants of behavior change and selecting appropriate strategies to address them. They can guide LLMs to generate more relevant and reliable responses.

Selection of behavior change theories. Following Moullin et al. [34]'s recommendations, our framework selection is guided by four key criteria: 1) purpose alignment with behavior change and implementation; 2) individual-level target capability; 3) capability to address implementation elements (e.g., barriers and strategies); and 4) contextual fit with individual lifestyles and chatbot interventions. After evaluating several frameworks, we choose Theoretical Domains Framework (TDF) [2] and the Behavior Change Technique (BCT) Taxonomy [32] because of their comprehensive coverage of behavior determinants and actionable techniques, and their suitability for individual-level interventions. Specifically, TDF offers 14 domains for identifying behavior change factors, such as knowledge, skills, belief about capabilities, and environmental context. BCT Taxonomy provides 93 specific techniques organized into 14 categories (e.g., goals and planning, feedback and monitoring) for implementing behavior change interventions. It offers practical methods for translating theoretical understanding into actionable interventions. This combination allows us to identify behavioral determinants and address them through actionable strategies.

We have also considered alternative frameworks. Behavior Change Wheel (BCW) [33] combines the COM-B model (Capability, Opportunity, Motivation—Behavior) with intervention functions and policy categories to design comprehensive behavior change interventions. It emphasizes policy enablers and broader intervention functions that are less directly applicable to our context. Similarly, we rule out Consolidated Framework for Implementation Research (CFIR) [8] since it highlights contextual factors such as organizational culture and policies. The Transtheoretical Model (TTM) [40] describes behavior change as a process through six stages, from precontemplation to termination. It focuses more on tailoring interventions for transitions of change over extended periods, which is less aligned with our immediate, adaptive interventions.

Consolidating theoretical frameworks. We first identify TDF domains that are most relevant for sleep behavior. Some domains like Social Role and Identity and Social Influences are excluded due to the challenges of simulating users' social connections in a chatbot context. Then, we select BCT categories that align with the identified TDF domains and are practical for a chatbot to deliver.

Finally, we consolidate the concepts derived from TDF and BCT into seven key techniques to inform our chatbot design. For example, the "intention" and "goals" domains from the TDF are merged as one category because they both represent a commitment to achieve a desired outcome. The TDF domain "Behavior Regulation" is merged with BCT "Feedback and Monitoring", both involving objective tracking and management of observed behaviors. BCT "Social Support" category is scoped to TDF "Emotion" domain. The final framework is detailed in Table 2.

3.4.2 Implementing Data-Driven and Theory-Guided Chatbot. We implement a multi-agent chatbot where each agent, powered by LLMs, performs specialized analytical tasks to generate comprehensive, personalized responses. They coordinate with each other by custom prompts and rules.

First, **Agent Coordinator** analyzes incoming messages and conversation history to determine which analytical tasks are needed. It then directs requests to corresponding specialized agents based on classification prompts. The agents then process these requests:

- **Technique Extraction Agent:** This agent identifies relevant behavior change strategies (see Table 2) for generating responses. It uses tailored prompts that fuse users' current messages, conversation history, and behavior change theories to instruct LLM selection of response strategies.
- **Personal Data Agent:** This agent retrieves and analyzes the user's personal data (Table 1), such as sleep patterns and activity levels, that are relevant to the user's messages. Given the user's chat messages and daily wearable data (organized as a dataframe) as input, the agent leverages the LLM's capabilities to generate Python code for data retrieval, transform, and analysis. Based on the code results and user message, another LLM reports the insights (e.g., average sleep duration over the past week).⁴
- **Recommendation Agent** is responsible for delivering context-aware personalized activity suggestions. Given the current users' context data (e.g., weather, time of day), our contextual MAB model suggests sleep-enhancing activities (e.g., run, walk). Then, to improve relevance and actionability, the agent prompts LLMs to adjust and tailor the model-generated activities based on the user's current contexts, locations, and situations implied in the messages. For example, it may suggest indoor treadmill running when it is raining outside or walking in the nearby park.

Finally, the **Response Agent** synthesizes outputs from all agents into cohesive replies that integrate data insights, personalized recommendations, and behavior change techniques. This modular approach decomposes complex chatbot's conversational and analytical reasoning tasks into specialized agents, ensuring responses are relevant, theory-grounded, and adaptive to user's specific circumstances.

4 SYSTEM IMPLEMENTATION

HEALTHGURU is implemented as an interactive web-based application using Streamlit, a Python framework to build data-driven

⁴If there are data and processing errors, the system will respond "I am sorry, I am not able to provide the information at the moment."

Table 2: Seven core techniques identified for sleep health improvement.

Technique Domain	Definition	Example
1. Consequences and Reinforcement	Discussing anticipated outcomes of sleep behaviors and providing feedback on user's actions	"Based on your Oura ring data, your consistent 10 PM bedtime has led to a 15% increase in your deep sleep. This improvement can enhance your memory and cognitive function."
2. Feedback and Monitoring	Tracking sleep patterns and providing users with insights into their progress	"I notice you've been using your Oura ring consistently. Let's review your sleep efficiency over the past week and identify areas for improvement."
3. Goals	Setting clear, achievable sleep-related objectives tailored to the user's current habits and desired outcomes	"Given your current average of 6 hours of sleep, shall we set a goal to gradually increase this to 7 hours over the next month?"
4. Knowledge	Providing users with tailored information about sleep health, addressing gaps in their understanding	"Did you know that exposure to blue light from devices before bedtime can disrupt your melatonin production? Let's discuss some strategies to minimize this effect."
5. Environmental Context and Resources	Addressing the user's physical sleep environment and available resources to optimize sleep conditions	"I see you live in a noisy urban area. Have you considered using a white noise machine to mask disruptive sounds during the night?"
6. Skills and Capabilities	Teaching users specific techniques to improve their sleep and building their confidence in implementing these strategies	"Let's practice a simple breathing exercise that can help you relax before bed. Inhale for 4 counts, hold for 7, and exhale for 8. How does that feel?"
7. Emotional Support	Providing empathy, encouragement, and motivation to support users' sleep improvement efforts	"I understand that changing sleep habits can be challenging. Remember, every small step you take is progress. How can I help you feel more confident about making these changes?"

web apps. Upon user content, the system accesses the user's location data via the browser, which is then used to retrieve local time, weather conditions, and temperature information through the WeatherAPI⁵. This contextual information is used to build the personalized activity recommendation model. Meanwhile, users' wearable data, including activity, sleep, and physiological data (e.g., heart rate), is accessed through Oura API⁶. The chatbot's core functionality is powered by OpenAI's language models. GPT-4o-mini is employed for managing general conversation flow, theoretical framework selection, and recommendation model integration and results refinement, ensuring a balance between response quality and speed. For the more complex task of retrieving and analyzing wearable data, we leverage GPT-4o to use the Python tool to generate analysis codes and results. Then, we build the LLM-based multi-agent framework and chatbot by using Python LangChain Library⁷. LangChain provides a suite of tools for constructing prompts, managing conversations, and orchestrating the flow between agents and LLMs via prompt chaining. It allows us to modularize the chatbot's functionalities into specialized agents that interact seamlessly. To enhance user experience, the system implements streamed output, reducing perceived latency and allowing for real-time interaction. The application is hosted on Amazon Elastic Compute Cloud (EC2), providing scalable computing capacity.

5 EVALUATION

We conducted an in-the-wild deployment study with 16 participants to evaluate our system compared to a baseline chatbot. We aimed to answer the following research questions:

- How does HEALTHGURU impact users' sleep health behaviors and outcomes?
- How effective are HEALTHGURU's activity recommendations and theory-guided conversations in motivating behavior change?
- What are users' perceived benefits and limitations towards using HEALTHGURU?

5.1 Study Design

We employed a within-subjects design to evaluate HEALTHGURU against a baseline system. The study lasted for eight weeks and consisted of three phases:

- **Observation (16 days):** Participants maintained their regular sleep and activity patterns, such that we can collect user historical data to initialize our models.
- **HEALTHGURU (20 days):** Our chatbot engaged users in theory-guided and data-driven conversations to motivate users' sleep-enhancing behavior change.
- **Baseline (20 days):** Participants used a ChatGPT-style chatbot capable of answering general health-related questions, providing general sleep advice, and answering questions about users' health data.

To mitigate ordering effects, we counterbalanced the Baseline and HEALTHGURU phases across participants.

5.2 System Implementation and Setup

Both HEALTHGURU and the Baseline were implemented as web-based applications hosted on AWS EC2, accessible via URL. Key differences between the two systems were:

⁵<https://www.weatherapi.com/>

⁶<https://cloud.ouraring.com/v2/docs>

⁷<https://github.com/langchain-ai/langchain>

- HEALTHGURU: Implemented as described in subsection 3.4, featuring theory-guided conversations and personalized activity recommendations.
- Baseline: Followed a similar framework but lacked the theory-guided approach and personalized activity recommendations of HEALTHGURU. Specifically, the baseline leverages agent coordinators to decide whether users ask questions about their data. If yes, the data insight agent uses a Python tool to retrieve and analyze the target data subsets and provide a response based on the conversation context. Otherwise, the chatbot responds to the user message directly.

Both systems were initialized with a system prompt to act as a sleep expert, helping users improve sleep quality.

5.3 Participants

We recruited (N=16) participants (age range: 18-44 years, 5 females, 10 males, 1 prefer not to say) from diverse sources, including email listservs and posters at university, personal connections, and social clubs (e.g., running clubs). Table 3 summarizes participant demographics. They were primarily Asians (9) and Whites (5); most held graduate degrees (14). Regarding work arrangements, 7 split time evenly between home and office, and 5 worked mainly in the office; almost all resided in urban areas.

Prior to the study, 7 participants did not own or wear a wearable device, and 9 reported using a wearable device for varying periods from less than 6 months to over 2 years. Participants were highly motivated—with 12 (75%) very or extremely motivated to improve physical activity and 13 (81.25%) to improve sleep health. Sleep patterns varied among participants, with 6 self-identifying as night owls, 5 as average sleepers, 3 as early birds, and 2 as irregular. The most common bedtimes were between 11 to 12 pm, and the most common wake times were between 6 to 9 am. This varied sample enabled the evaluation of our sleep health chatbot across different profiles, work arrangements, and sleep habits.

5.4 Study Procedures

5.4.1 Pre-study Preparations. We took comprehensive measures to ensure ethical conduct and participant readiness. First, participants were thoroughly informed about the study protocols, such as data collection methods and privacy measures. They signed consent forms detailing the study’s purpose, procedures, risks, and benefits. Participants were instructed on how to interact with the chatbots, asking questions about their sleep and activity data and requesting suggestions. Sample questions were provided to help initiate conversations with both systems. For data protection, participants were assigned unique identifiers (user ID and password) to access the chatbots. They were informed about the chatbots’ access to their Oura ring metrics and instructed not to input sensitive content. Oura rings were provided along with instructions for granting and later revoking API access for data collection.

5.4.2 Study Phases. After the observation periods, participants used either the Baseline or HEALTHGURU for three weeks. Finally, semi-structured interviews (45 minutes) were conducted at the study’s end to gather in-depth feedback on both systems. Throughout the study, participants were asked to wear their Oura ring

Table 3: Participant demographics.

Characteristic	Category	N (%)
Age	18 to 24	3 (18.75)
	25 to 34	9 (56.25)
	35 to 44	4 (25.00)
Gender	Female	5 (31.25)
	Male	10 (62.50)
	Prefer not to say	1 (6.25)
Race/Ethnicity	Asian	9 (56.25)
	Black or African American	1 (6.25)
	Other	1 (6.25)
	White	5 (31.25)
Occupation	Full-time employed	8 (50.00)
	Full-time employed and Student	2 (12.50)
	Not employed	1 (6.25)
	Student	5 (31.25)
Wearable Device Usage	No prior ownership	6 (37.50)
	Own but never use	1 (6.25)
	Less than 6 months	3 (18.75)
	6 months to 2 years	3 (18.75)
	Over 2 years	3 (18.75)
Motivation to Improve Physical Activity	Extremely motivated	4 (25.00)
	Very motivated	8 (50.00)
	Somewhat motivated	4 (25.00)
Motivation to Improve Sleep Health	Extremely motivated	5 (31.25)
	Very motivated	8 (50.00)
	Somewhat motivated	2 (12.50)
	Not very motivated	1 (6.25)
Sleep Pattern	Night owl	6 (37.50)
	Average sleeper	5 (31.25)
	Early bird	3 (18.75)
	Irregular	2 (12.50)

continuously and carry their iPhones, complete daily questionnaires, participate in post-system questionnaires, and engage in exit interviews at the study’s conclusion.

Daily questionnaires collected users’ perceived sleep quality, motivation levels, and responses to chatbot recommendations. The daily survey included the following key components: (1) sleep quality rating (5-point scale: “Terrible” to “Excellent”); (2) motivation level to improve sleep/activity (5-point scale); (3) reasons for low motivation (if applicable); (4) relevance of chatbot recommendations (5-point scale: “Not Relevant at all” to “Very Relevant”); (5) adherence to recommendations (Yes/No);

Post-system surveys assessed user experience on a 5-point Likert scale: (1) technical aspects (e.g., response time, accuracy, user interface design, overall usability); (2) personalization and relevance of recommendations; (3) impact on understanding sleep-activity relationships; (4) motivation and behavior change; (5) perceived improvements in sleep quality. Open-ended questions were also collected on behavior changes and system strengths and limitations.

Exit interviews gathered feedback on (1) overall impressions of both systems; (2) perceived impacts on sleep patterns and physical activity; (3) perceived engagement and usage over time; (3) most and least helpful features; (4) challenges in adopting recommendations; (5) suggestions for improvement; (6) likelihood of continued use.

5.5 Data Collection

We collected the following data daily throughout the study.

Wearable data. Sleep and activity data were collected using the Oura ring (see Table 1). The metrics included duration, efficiency, quality, average HRV, and lowest heart rate, stress levels, and readiness scores. Activity types and scores were also recorded.

Questionnaires and interview data. We collected daily questionnaires, post-system questionnaires, and exit interview feedback (detailed in subsection 5.1).

System usage data. We collected data on users' conversation history with the chatbots. Based on that, we computed the frequency of interactions and length of conversations.

5.6 Data Security and Ethical Considerations

All data was securely stored on a university-hosted Amazon Web Services (AWS) server. To protect participants' privacy, personal information was stored separately from research data, with identifying details securely kept on the EZ backup system. To ensure ethical use of chatbots, We implemented guardrails using OpenAI's moderation API⁸ to detect and block the generation of potentially harmful content. Regular check-ins were also conducted to address participants' concerns about chatbot responses.

5.7 Data analysis

We conducted analyses corresponding to each research question:

5.7.1 Impact on Sleep Health Behaviors and Outcomes.

- Paired t-tests to compare sleep and activity metrics (e.g., sleep duration, sleep efficiency, and readiness score) between Baseline and HEALTHGURU conditions
- Descriptive statistics (mean and standard deviation) for stress level distribution comparison
- Thematic analysis of user responses regarding perceived impacts on sleep patterns and physical activity

5.7.2 Effectiveness of Recommendations and Theory-guided Conversations.

- Wilcoxon signed-rank test to compare personalization and relevance ratings, and adherence rates between Baseline and HEALTHGURU conditions
- Descriptive statistics for the frequency of behavior change techniques employed by HEALTHGURU
- Thematic analysis of user feedback on the most and least helpful features of HEALTHGURU
- Content analysis of chat logs to identify prevalent behavior change techniques and their implementation
- Expert feedback on system responses to users' messages

5.7.3 User Engagement and Perceived Benefits and Limitations.

- Paired t-tests to compare user engagement metrics (ratio of active days, conversation length) between Baseline and HEALTHGURU conditions
- Linear regression analysis to examine engagement trends over time for both systems

- Wilcoxon signed-rank test to compare user ratings on various aspects of system performance, usability, and perceived benefits between Baseline and HEALTHGURU conditions
- Thematic analysis of interview responses and open-ended survey questions to identify:
 - Perceived benefits and limitations of HEALTHGURU
 - Challenges in adopting recommendations
 - Suggestions for improvement
- Descriptive statistics of user satisfaction ratings

6 RESULTS AND ANALYSIS

We analyzed quantitative data from user interactions, Oura ring metrics, ratings, and qualitative participant feedback to provide a comprehensive comparison of HEALTHGURU and the Baseline.

Table 4: Comparative analysis of sleep and activity metrics collected by Oura for Baseline and HEALTHGURU.

Metric	Baseline		HEALTHGURU		P-val	Sig.?
	Mean	SD	Mean	SD		
Sleep Duration (hrs)	6.48	0.59	6.84	0.63	0.0268	Yes
Sleep Efficiency (%)	80.27	5.20	80.58	4.21	0.8490	No
Average Breath	15.49	1.17	15.50	1.10	0.5882	No
Average HRV	53.45	19.14	49.15	19.85	0.6200	No
Lowest Heart Rate	56.99	7.64	57.21	8.32	0.9537	No
Sleep Score	75.28	13.23	78.34	13.77	0.4764	No
Activity Score	71.37	14.18	75.41	13.94	0.0268	Yes
Readiness Score	76.44	6.34	77.91	6.29	0.4570	No

Table 5: Post-system questionnaire ratings of Baseline and HEALTHGURU. Medians and interquartile ranges (IQR) are reported. Wilcoxon signed-rank test was used for statistical significance testing.

Aspect	Baseline	HEALTHGURU	W	p-val	Sig?
	Med (IQR)	Med (IQR)			
Speed of chatbot	4.00 (2.00)	4.00 (1.00)	5.00	0.4922	No
Accuracy of responses	2.50 (2.25)	4.00 (0.00)	20.00	0.0375	Yes
User interface design	4.00 (1.00)	4.00 (1.25)	42.00	0.7815	No
Ease of accessing data	4.00 (1.00)	4.00 (1.00)	7.50	1.0000	No
Overall system usability	3.50 (1.25)	4.00 (0.00)	37.50	0.2850	No
Personalized recommendations	4.00 (2.00)	4.00 (0.00)	9.00	0.0308	Yes
Feasible recommendations	3.50 (1.25)	4.00 (1.00)	6.00	0.1597	No
Understood activity-sleep relation	3.50 (2.00)	4.00 (2.00)	8.00	0.0386	Yes
Learned new strategies	4.00 (1.00)	4.00 (1.00)	33.00	1.0000	No
Effective use of Oura data	3.00 (2.00)	4.00 (1.00)	4.00	0.0126	Yes
Increased motivation	3.50 (1.25)	4.00 (1.25)	8.00	0.0126	Yes
Made changes to routine	4.00 (1.50)	3.50 (1.00)	16.00	0.4170	No
Noticed sleep improvements	3.00 (1.25)	4.00 (1.00)	11.00	0.0187	Yes
Understood daily impact on sleep	2.50 (2.00)	3.00 (1.00)	2.50	0.0461	Yes
Increased awareness of patterns	4.00 (2.50)	4.00 (2.00)	20.00	0.1177	No

⁸<https://platform.openai.com/docs/guides/moderation>

6.1 Impact on Sleep Health Activities and Outcomes

6.1.1 Sleep Metrics. Table 4 shows that participants using HEALTHGURU had a significant increase in sleep duration—an average improvement of 22 minutes per night ($t = -2.59, p < 0.05$). However, analysis of sleep scores revealed no statistically significant difference ($p = 0.47$) between HEALTHGURU (78.34) and the Baseline (75.28). Similarly, while sleep efficiency showed a slight increase (80.27% vs. 80.58%) and HRV exhibited a minor decrease (53.45 vs. 49.15), these differences were not statistically significant. The average breath rate and lowest heart rate during sleep remained relatively constant between the two phases.

6.1.2 Activity and Readiness. HEALTHGURU significantly increased physical activity scores, improving from 71.37 (SD = 14.18) to 75.41 (SD = 13.94) ($t = -2.59, p < 0.05$), indicating its ability to motivate participants to engage in more activity. Although readiness scores improved slightly (76.44 to 77.91), this change was not significant ($t = 0.77, p = 0.457$).

6.2 Effectiveness of Recommendations and Behavior Change Techniques

6.2.1 Personalization, Relevance, and Adherence of Recommendations. As shown in Table 5, HEALTHGURU significantly outperformed the Baseline in providing personalized recommendations, confirmed by the Wilcoxon signed-rank test ($W=9.00, p=0.0308$). Participants rated HEALTHGURU’s advice as more relevant to their needs, with significantly higher relevance scores (Median (IQR): 3.81 (0.34) vs. 3.34 (0.35); $W = 20.00, p = 0.023$). Qualitative responses further highlighted HEALTHGURU’s improved contextualization, with one user stating, “I think having the chatbot know the current weather and temperature outside is motivating. It recommended me to go out for a walk in a sunny afternoon, which I sometimes did.” However, participants felt that the Baseline’s suggestions were “boring” and “generic” which they already knew, such as maintaining 7 hours of sleep, watching coffee intake, etc.

The HEALTHGURU’s ability to provide location-context-aware recommendations was particularly appreciated. One user reported, “The chatbot suggested ‘a light jog around Four Freedoms Park in Roosevelt Island in the later afternoon, considering the sunny weather and cooling down the temperature.’ It felt like the advice was tailored just for me and my surroundings.” Some participants noticed that HEALTHGURU knew how to adapt their daily routines to the environment. One user said, “I remember it was a bit rainy and cold that morning, the chatbot suggested indoor activities like yoga and treadmill rather than outdoor running, which I normally do.” Another added, “Once, it recommended indoor strength exercises with air conditioning due to the high temperature. It felt very considerate.”

HEALTHGURU also demonstrated flexibility in adapting recommendations to users’ specific situations. A user mentioned that one time the chatbot recommended an outdoor walk, but he was recovering from COVID. He told the chatbot and the chatbot put forward 15-30 minutes of gentle stretching and yoga with some poses to try. Another one appraised, “When I mentioned I don’t like meditation, the chatbot quickly switched the suggestions to mindful breathing exercises instead. It felt like it really listened to my preferences.”

This enhanced personalization resulted in higher adherence rates. Participants followed through on a median of 47.94% (IQR: 34.27%) of HEALTHGURU’s recommendations, which is significantly higher compared to a median of 36.71% (IQR: 15.49%) for the Baseline ($W = 28.00, p = 0.0386$). Users frequently mentioned the relevance and practicality of the recommendations as key factors in their increased adherence. However, there was no significant difference between the two systems regarding the perceived feasibility of the recommendations.




Figure 3: Comparison of behavior change strategies covered in Baseline and HEALTHGURU systems.

6.2.2 Implementation of Behavior Change Strategies. Post-system questionnaire analysis results (in Figure 3) reveal that HEALTHGURU covered a broader range of behavior change strategies compared to the Baseline system. Notably, HEALTHGURU showed large improvements in providing feedback on sleep patterns and physical activity (15 vs. 10 participants), addressing environmental context (12 vs. 8), enhancing skills and capabilities for better sleep (7 vs. 3), providing emotional support (10 vs. 3), and facilitating goal-setting (8 vs. 2). Both systems performed similarly in addressing consequences and reinforcement (9 participants each).

Analysis of user conversations processed by our multi-agent framework identified seven technique domains, with knowledge being most prevalent (23.8% of interactions), followed by feedback and monitoring (17.9%), goals (17.6%), skills and capabilities (16.9%), emotion support (9.4%), consequences and reinforcement (8.2%), and environment context and resources (6.2%). Each strategy was exemplified in HEALTHGURU’s responses. For instance, HEALTHGURU provided knowledge about potential causes of sleep disturbances and actionable advice:

User: “What are the possible reasons for waking up at midnight?”

HEALTHGURU: “Waking up at midnight can often be linked to stress, sleep disorders, or environmental factors like noise and temperature.”

HEALTHGURU also encouraged ongoing monitoring and maintaining of positive sleep patterns.

HEALTHGURU: “Your sleep efficiency has been consistently high over the past few days, averaging around 91%. This indicates that you’re spending a significant portion of your time in bed actually sleeping, which is a positive sign for your overall sleep quality. Keep monitoring this trend to maintain or improve your sleep health.”

In addition, HEALTHGURU provided advice to enhance users’ skills for improving sleep quality.

HEALTHGURU: “Practice a relaxing pre-sleep routine, such as reading or gentle stretching, and incorporate mindfulness or breathing exercises to lower stress levels before bedtime.”

Furthermore, qualitative feedback supported the effectiveness of HEALTHGURU’s behavior change strategies. Users appreciated the emotional support and personalized feedback. One stated, “That’s great! I’m proud of you for being open to trying.” Another added, “Feedback on my sleep patterns related to activity... I am trying to be more active. Actually, I just finished a workout. That’s been great!”

6.2.3 Expert Feedback. Three experts (E1-E3) evaluated system responses to users’ messages regarding scientific validity, data accuracy, and applicability. E1 is a health psychologist and the director of a research center for health promotion by physical activity. E2 is an assistant professor in sports and exercise science. E3 is a PhD researcher who has extensive experience (e.g., publications) in sleep and behavior sensing. We sampled 12 anonymized conversations from study participants using HEALTHGURU, along with corresponding wearable data. The samples covered the system’s three key capabilities: behavior change techniques, data insights, and activity recommendations, with at least four examples of each⁹.

Experts found that HEALTHGURU could accurately retrieve and analyze the relevant wearable data according to users’ questions. For example, when asking about “sleep last night”, HEALTHGURU could locate the relevant data attributes, such as sleep duration, efficiency, HRV, and lowest heart rate, on the requested date. Moreover, E3 added that it was “user-friendly” to convert original sleep duration units from seconds to hours. Experts generally thought the responses were aligned with established sleep and behavior change theories. They also appreciated the actionability of HEALTHGURU’s activity recommendations, such as “take a route that passes a nearby park” while walking to grocery shops and “walk around the terminal (while at the airport)”.

However, they pointed out areas for improvement. E3 required a more in-depth scientific interpretation of the users’ wearable data (“What an HRV value of 64 means for the user’s sleep”) to facilitate users’ understanding beyond simple data summarization. Some proposed solutions to sleep problems were deemed general and less useful in motivating users to act (“15-30 min yoga”). More context information (about “waking up at midnight” and “broken heart”) might be needed to infer users’ psychological and physical needs and connect them with appropriate behavior change techniques (E1). E2 noted that certain conclusions felt less convincing, with “only one data point used as evidence (in the response)”.

⁹Individual examples could span multiple categories. They are included in Supplementary Material.

Table 6: Comparison of user engagement metrics between Baseline and HEALTHGURU systems.

Metric	Baseline	HEALTHGURU	t-stat	p-val
	Mean (SD)	Mean (SD)		
User Engagement	0.31 (0.19)	0.39 (0.15)	-2.47	0.026*
Conversation Length	4.84 (2.16)	6.83 (2.67)	-2.18	0.045*

* Indicates statistical significance at $p < 0.05$.




Figure 4: Comparison of user engagement over time using Baseline and HEALTHGURU systems.

6.3 User Engagement Analysis

To evaluate the effectiveness of HEALTHGURU compared to the Baseline system, we analyzed key metrics of user engagement (Table 6): the ratio of active days to total phase days and conversation length. These metrics provided insights into how frequently and deeply users interacted with each system over the course of the study.

6.3.1 Engagement Metrics. User engagement, defined as the ratio of active days to total phase days, showed a significant improvement with HEALTHGURU. The average engagement ratio increased from 0.31 (SD = 0.19) with the Baseline system to 0.39 (SD = 0.15) with HEALTHGURU, a statistically significant improvement ($t = -2.47$, $p = 0.026$). This result indicates that users interacted with HEALTHGURU more consistently over the study period, suggesting enhanced long-term engagement. The average number of messages in each conversation also increased significantly from 4.84 (SD = 2.16) in the Baseline system to 6.83 (SD = 2.67) with HEALTHGURU ($t = -2.18$, $p = 0.045$). This suggests that users engaged in longer conversations with HEALTHGURU. The increased conversation length may reflect users’ higher levels of engagement and HEALTHGURU’s ability to provide more comprehensive and tailored responses, which encouraged users to explore health topics in greater detail.

6.3.2 Engagement Trends Over Time. To understand how user engagement evolved throughout the study, we analyzed daily active user counts for both systems over the study period (in Figure 4). The Baseline system showed a significant downward trend in engagement with linear regression analysis (slope = -0.018 , $p = 0.0025$, $R^2 = 0.41$). It started with 10 active users, which decreased sharply after the first few days. In contrast, HEALTHGURU demonstrated a

more stable engagement pattern. While it started with similar initial engagement (11 active users), it maintained a relatively consistent level of user activity throughout the study period, with periodic fluctuations between 4-9 users. HEALTHGURU's trend analysis showed a much gentler downward slope (-0.0037) that was not statistically significant ($p = 0.49$, $R^2 = 0.027$), indicating a better sustained user interest over time.

The combination of higher engagement ratios, longer conversations, and more stable daily active users demonstrate HEALTHGURU's effectiveness in maintaining sustained user engagement. This is particularly critical for the long-term success of sleep health interventions, as consistent engagement offers users more opportunities to receive personalized advice, reflect on their sleep patterns, and implement meaningful changes to their routines.

6.4 System Performance and Usability

We compared HEALTHGURU with the Baseline system in terms of system response, user interface design, and usability metrics (Table 5). Our analysis reveal significant improvements in the HEALTHGURU system compared to the Baseline system in terms of response accuracy. The median accuracy ratings of responses improved from 2.50 (IQR = 2.25) to 4.00 (IQR = 1.00), $W = 20.00$, $p = 0.038$. However, there were no significant differences in speed, user interface design, ease of accessing data, or overall system usability between the two systems. This suggests that while HEALTHGURU maintained the user-friendly aspects of the Baseline system, it significantly enhanced the quality and relevance of its responses. Qualitative feedback supported this finding. For instance, one user noted, *"The chatbot gave very good advice about my sleep patterns, such as noticing my lack of sleep or lack of exercise. It gave great recommendations personalized to my surroundings and fitness."*

6.5 Users' Perceived Benefits and Limitations

We analyzed user feedback on the perceived benefits and limitations of HEALTHGURU, focusing on aspects such as personalization, usefulness, and engagement.

6.5.1 Perceived Benefits. HEALTHGURU demonstrated significant improvements in several key areas compared to the Baseline system.

Participants reported a significantly better understanding of sleep-activity relationships when using HEALTHGURU. The median score for this aspect increased from 3.50 (IQR: 2.00) to 4.00 (IQR: 2.00) ($W = 8.00$, $p = 0.039$). Similarly, users' understanding of how daily activities impact sleep also improved significantly, with median scores rising from 2.50 (IQR: 2.00) to 3.00 (IQR: 1.00) ($W = 2.50$, $p = 0.046$). One user remarked, *"After it told me my sleepiness is due to inactivity rather than bad sleep, I am trying to be more active."*

HEALTHGURU significantly outperformed the Baseline system in the effective use of Oura ring data. The median score increased from 3.00 (IQR: 2.00) to 4.00 (IQR: 1.00) ($W = 4.00$, $p = 0.013$). This suggests that HEALTHGURU was more effective in leveraging user data to provide meaningful insights.

HEALTHGURU also had a notable positive impact on motivation and behavior change. The median score for increased motivation to improve sleep habits rose from 3.50 (IQR: 1.25) to 4.00 (IQR: 1.25) ($W = 8.00$, $p = 0.013$). Additionally, users reported greater sleep

improvements after following its advice, with scores increasing from 3.00 (IQR: 1.25) to 4.00 (IQR: 1.00) ($W = 11.00$, $p = 0.019$).

Both systems were effective in increasing users' awareness of their sleep patterns (median = 4.00 for both), but HEALTHGURU seemed to inspire more specific behavioral changes. One user reported, *"Talking to the chatbot definitely made me more aware of my sleep. Also, there were times I've been paranoid about not sleeping enough, but chatbot told me my sleep score is fine, most likely I am sleepy because of my (low) activity score."*

6.5.2 Perceived Limitations and User Suggestions. Despite the overall positive feedback, users identified several limitations and areas for improvement in HEALTHGURU, particularly in data accuracy, actionability of insights, data integration, and user experience.

Some participants reported data accuracy issues in Oura ring. One user noted the sleep score did not align well with their subjective experience. Some mentioned that sometimes the activities were not captured correctly or not detected at all. One said, *"The app showed that I had a nap, but I was just sitting on chair playing my phone."* Moreover, some participants felt the insights could be more actionable and comprehensive. As one user put it, *"The chatbot told me stress is likely a cause, but it didn't give actionable insights. It suggested meditation, but I still had to find my own tutorials."* Participants desired system capability to integrate and track factors like daily diet and routines that could influence sleep.

Several usability concerns were raised, including the need to switch between multiple apps, the absence of effective data visualization, and a preference for voice commands over typing. To address these issues, participants suggested developing an integrated platform that seamlessly combines chatbot interactions and data access within a single interface. In addition, users recommended implementing a dedicated data visualization dashboard to summarize key metrics like stress levels, sleep quality, and their trends over time, making the information easier to interpret. Many participants emphasized the value of voice-based interactions to facilitate system usage, particularly when on the go or in situations where typing is inconvenient. Participants also highlighted the need of proactive engagement features, such as push notifications to highlight progress, track goals, and provide positive reinforcement.

7 DISCUSSION

We introduced HEALTHGURU, an LLM-powered chatbot that delivers data-driven, theory-guided sleep health support through a multi-agent framework. This framework features a contextual multi-armed bandit model for adaptive activity recommendations with behavior change strategies in conversational flows.

7.1 Effectiveness of Data-Driven, Theory-Guided Approach

An eight-week deployment study with 16 participants demonstrated that HEALTHGURU substantially outperformed a baseline system in promoting healthier sleep behaviors. By integrating wearable data (e.g., Oura ring metrics) with contextual factors such as time, weather, and location, HEALTHGURU delivered recommendations that participants consistently found more personally relevant, context-aware, and actionable. This improved personalization and

contextual relevance led to notably higher adherence rates, a 22-minute increase in average sleep duration, and improved activity scores. Participants were more inclined to follow HEALTHGURU's advice—such as taking a light jog in a nearby park—than those provided by the baseline. Moreover, the inclusion of evidence-based behavior change strategies led to a better understanding of sleep-activity relationships and increased motivation for healthy habits.

7.2 Balancing Automated Personalization and User Agency

While participants valued HEALTHGURU's personalized capabilities, the study raises important questions about the balance between automated, AI-driven recommendations and user autonomy. Some participants expressed the need to maintain control over their health decisions. For example, one mentioned, *“When it suggests going for a walk because it's sunny, but I feel dehydrated and would prefer to rest.”* This highlights the importance of allowing users to override or adapt system recommendations based on their immediate feelings or circumstances. Currently, HEALTHGURU supports user agency through its conversational interface, enabling users to ask for clarifications or alternatives. Future approaches can consider interaction designs that allow users to easily specify what aspects of information at what levels of detail should be presented to empower them to make informed decisions.

Moreover, the persuasive nature of personalized recommendations, while effective for behavior change, raises ethical questions about autonomy and informed consent. As one participant noted, *“I find myself following the chatbot's advice more often than not.”* This underscores the need for transparent communication about the system's capabilities and limitations and the importance of framing AI recommendations as suggestions rather than directives. Future systems should consider implementing more explicit user control features, such as preference settings for recommendation frequency or type, and clear options for users to provide feedback on or override suggestions. Additionally, incorporating periodic reminders of the system's AI nature and encouraging users to consult healthcare professionals for serious concerns could help maintain an appropriate balance between automated support and user agency.

7.3 Challenges in Sustained Engagement and Behavior Change

Our study revealed significant challenges in maintaining long-term engagement and translating increased awareness into consistent behavior change.

- **Declining novelty:** Participants mentioned that initial curiosity and interest in the chatbot's capabilities decreased over time. As one user noted, *“At first, I was curious about what it could suggest, but over time, I felt less motivated to check in with the chatbot.”*
- **Cognitive load and system interaction barriers:** Manual text input introduced a cognitive burden, as evidenced by user feedback requesting voice interactions. Some highlighted the need for better support for handling recurring queries and entry patterns to simplify engagement.

- **Recommendation fatigue:** As participants became familiar with the system's suggestion patterns, they reported decreasing perceived value from recommendations as they began to feel repetitive to users.
- **Changing life circumstances and competing priorities:** Real-life situations, such as travel and recovering from illness, hindered users' ability to engage with the chatbot. *“I knew my sleep was pretty bad lately, but I had to sacrifice my sleep since there is an urgent deadline ahead of me.”*

Based on user study results, future systems can consider the following designs. 1) To sustain novelty and interest, systems can introduce adaptive challenge levels that involve users' progress toward their goals, similar to techniques used in gamification. They can also integrate periodic novel content and social features. 2) To improve system interactions, multi-modal interactions (e.g. voice commands) and data visualizations should be enabled to enhance natural and convenient data interactions and analytics. Implementation of smart defaults and interaction shortcuts for common queries could simplify engagement. 3) To mitigate recommendation fatigue, systems can continuously expand contextual factors (e.g., seasonal shift) into recommendation logic. Recommendation timing can be optimized via intelligent notification scheduling based on user receptivity patterns. 4) Systems can support customizable interaction levels, allowing users to adjust the intensity of chatbot engagement based on their current capacity and preferences.

7.4 Implications for Designing AI-Powered Health Support Systems

Our study findings suggest three design considerations for future AI-powered health support chatbots.

Emphasizing transparency and reliability to build trust.

The ability of HEALTHGURU to explain the rationale behind its suggestions was highly valued by participants and experts. This can foster trust and motivate adherence to recommendations. By linking advice to personal data and contextual factors, the chatbot helped users understand the benefits of suggested actions. Future health chatbots should prioritize explainability and transparency, clearly communicating how recommendations are generated and how user data is utilized. This approach aligns with our implementation goals in HEALTHGURU and is essential for building long-term user trust.

Applying theory-guided algorithms. Operationalizing behavior change theories within data-driven AI systems, as implemented in HEALTHGURU, improved response quality and intervention outcomes. The chatbot's use of the TDF and BCT enabled more effective support for behavior change. However, careful implementation is required to ensure that theoretical models are appropriately adapted to the AI context. Future systems should continue to incorporate theoretical guidance to enhance data-driven AI methods.

Integration with wearables and health systems. Our study reveals opportunities for deeper integration of chatbots with wearables and health systems.

- **Seamless user experience:** Systems should use a unified interface that combines wearable data visualization, chatbot interactions, and health insights. This will address user frustrations with context switching and enhance engagement.

- Adaptive real-time data analytics: Enhancing HEALTHGURU's integration of Oura Ring data, future designs can provide instant, personalized feedback through dynamic dashboards and predictive analytics (e.g., "Reducing screen time tonight may improve deep sleep by 10%"). These features empower users to make more timely, informed health decisions.
- Complementing existing health systems: AI chatbots should interpret data from various wearables, providing actionable advice that bridges the gap between users and their health information. This approach can supplement traditional health-care by supporting daily health management and enabling personalized, continuous interventions.

7.5 Limitations and Future Work

While HEALTHGURU demonstrated promising results in promoting sleep health, our study has several limitations.

- **Data accuracy challenges:** Our system primarily relied on Oura ring, which, while comprehensive and accurate, may still contain missing records, errors, and uncertainty that may impact system performance. Currently, if the system encounters data processing errors, HEALTHGURU will respond "I am sorry, I am not able to provide the information at the moment." In addition, users can cross-check raw data tables to judge data quality and system responses. In the future, systems should communicate data quality to users and give more weight to high-confidence data points when generating responses.
- **Contextual factors:** While we incorporated time, location, temperature, and weather into our recommendation agent, other relevant factors such as social context, diseases and medications, or work schedules were not included due to technical and privacy constraints. This limitation may affect the precision of personalized recommendations.
- **Model reliability evaluation and enhancement:** To improve model reliability, our system adopted multi-agent designs that decompose conversational and analytical tasks into specialized agents to form the final response. The system's reliability was further evaluated through both user and expert judgment. We can scale up the evaluation by developing automated fact-checking methods using external knowledge bases [21, 43]. Those data resources can also enhance the LLM reliability via retrieval-augmented generation and finetuning.
- **Study sample size and diversity:** Our current user study involved 16 participants (primarily young adults). Expanding the demographic diversity and sample size can improve the generalizability of findings.
- **Study duration:** Longer deployment studies beyond the current eight weeks are needed to assess sustained behavior change and long-term engagement.

8 CONCLUSIONS

We presented HEALTHGURU, a novel sleep health chatbot that integrates data-driven insights, behavior change theories, and adaptive recommendations via a multi-agent LLM framework. Integrating a

contextual multi-arm bandit model, HEALTHGURU delivers personalized activity recommendations and theory-grounded conversations to support behavior change. Our eight-week deployment study demonstrated significant advantages over a baseline system, with participants showing improved sleep duration, more consistent engagement, and increased motivation to adopt healthy sleep habits. However, our study also revealed challenges in maintaining long-term engagement and supporting sustained behavior change. Future work can improve data visualizations, implement more proactive features, and personalize recommendations to broader individual contexts and preferences.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their feedback. This work was supported by grants from National Science Foundation (2212351, 2212175) and National Institutes of Health (R01AG080991, R01AG076234). Fei Wang is the corresponding author.

REFERENCES

- [1] Mahyar Abbasian, Iman Azimi, Amir M Rahmani, and Ramesh Jain. 2023. Conversational Health Agents: A Personalized LLM-Powered Agent Framework.
- [2] Lou Atkins, Jill Francis, Rafat Islam, Denise O'Connor, Andrea Patey, Noah Ivers, Robbie Foy, Eilidh M Duncan, Heather Colquhoun, Jeremy M Grimshaw, et al. 2017. A Guide to Using the Theoretical Domains Framework of Behaviour Change to Investigate Implementation Problems. *Implementation Science* 12 (2017), 1–18.
- [3] A Bandura. 1977. Self-Efficacy: Toward a Unifying Theory of Behavioral Change. *Psychological Review* (1977).
- [4] Jared S Bauer, Sunny Consolvo, Benjamin Greenstein, Jonathan Schooler, Eric Wu, Nathaniel F Watson, and Julie Kientz. 2012. ShutEye: Encouraging Awareness of Healthy Sleep Recommendations with a Mobile, Peripheral Display. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM New York, NY, USA, 1401–1410.
- [5] Timothy W Bickmore and Rosalind W Picard. 2005. Establishing and Maintaining Long-Term Human-Computer Relationships. *ACM Transactions on Computer-Human Interaction (TOCHI)* 12, 2 (2005), 293–327.
- [6] Eun Kyoung Choe, Bongshin Lee, Matthew Kay, Wanda Pratt, and Julie A Kientz. 2015. SleepTight: Low-Burden, Self-Monitoring Technology for Capturing and Reflecting on Sleep Behaviors. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 121–132.
- [7] Justin Cosentino, Anastasiya Belyaeva, Xin Liu, Nicholas A Furlotte, Zhun Yang, Chace Lee, Erik Schenck, Yojan Patel, Jian Cui, Logan Douglas Schneider, et al. 2024. Towards a Personal Health Large Language Model. *arXiv preprint arXiv:2406.06474* (2024).
- [8] Laura J Damschroder, David C Aron, Rosalind E Keith, Susan R Kirsh, Jeffery A Alexander, and Julie C Lowery. 2009. Fostering Implementation of Health Services Research Findings into Practice: A Consolidated Framework for Advancing Implementation Science. *Implementation Science* 4 (2009), 1–15.
- [9] Nediya Daskalova, Bongshin Lee, Jeff Huang, Chester Ni, and Jessica Lundin. 2018. Investigating the Effectiveness of Cohort-Based Sleep Recommendations. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 3 (2018), 1–19.
- [10] Nediya Daskalova, Danaë Metaxa-Kakavouli, Adrienne Tran, Nicole Nugent, Julie Boergers, John McGeary, and Jeff Huang. 2016. SleepCoach: A Personalized Automated Self-Experimentation System for Sleep Recommendations. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, 347–358.
- [11] Zachary Englhardt, Chengqian Ma, Margaret E. Morris, Chun-Cheng Chang, Xuhai "Orson" Xu, Lianhui Qin, Daniel McDuff, Xin Liu, Shwetak Patel, and Vikram Iyer. 2024. From Classification to Clinical Insights: Towards Analyzing and Reasoning About Mobile and Behavioral Health Data With Large Language Models. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 8, 2, Article 56 (May 2024), 25 pages. <https://doi.org/10.1145/3659604>
- [12] Kathleen Kara Fitzpatrick, Alison Darcy, and Molly Vierhile. 2017. Delivering Cognitive Behavior Therapy to Young Adults with Symptoms of Depression and Anxiety Using a Fully Automated Conversational Agent (Woebot): A Randomized Controlled Trial. *JMIR Mental Health* 4, 2 (2017), e7785.
- [13] Billie Giles-Corti and Robert J Donovan. 2002. The Relative Influence of Individual, Social and Physical Environment Determinants of Physical Activity. *Social Science & Medicine* 54, 12 (2002), 1793–1812.

- [14] Wendy Hardeman, Julie Houghton, Kathleen Lane, Andy Jones, and Felix Naughton. 2019. A Systematic Review of Just-in-Time Adaptive Interventions (JIT-Is) to Promote Physical Activity. *International Journal of Behavioral Nutrition and Physical Activity* 16 (2019), 1–21.
- [15] Florian Herbolzheimer, Laura A Schaap, Mark H Edwards, Stefania Maggi, Ángel Otero, Erik J Timmermans, Michael D Denkiner, Suzan van der Pas, Joost Dekker, Cyrus Cooper, et al. 2016. Physical Activity Patterns among Older Adults With and Without Knee Osteoarthritis in Six European Countries. *Arthritis Care & Research* 68, 2 (2016), 228–236.
- [16] Matthew Kay, Eun Kyoung Choe, Jesse Shepherd, Benjamin Greenstein, Nathaniel Watson, Sunny Consolvo, and Julie A Kientz. 2012. Lullaby: A Capture & Access System for Understanding the Sleep Environment. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*. 226–234.
- [17] Michael P Kelly and Mary Barker. 2016. Why Is Changing Health-Related Behaviour So Difficult? *Public Health* 136 (2016), 109–116.
- [18] Young-Ho Kim, Sungdong Kim, Minsuk Chang, and Sang-Woo Lee. 2022. Leveraging Pre-Trained Language Models to Streamline Natural Language Interaction for Self-Tracking. *arXiv preprint arXiv:2205.15503* (2022).
- [19] Jochen Klenk, Gisela Büchele, Kilian Rapp, Sebastian Franke, Richard Peter, ActiFE Study Group, et al. 2012. Walking on Sunshine: Effect of Weather Conditions on Physical Activity in Older People. *J Epidemiol Community Health* 66, 5 (2012), 474–476.
- [20] Ahmet Baki Kocaballi, Shlomo Berkovsky, Juan C Quiroz, Liliana Laranjo, Huong Ly Tong, Dana Rezagadegan, Agustina Briatore, and Enrico Coiera. 2019. The Personalization of Conversational Agents in Health Care: Systematic Review. *Journal of Medical Internet Research* 21, 11 (2019), e15360.
- [21] Neema Kotonya and Francesca Toni. 2020. Explainable Automated Fact-Checking for Public Health Claims. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (Eds.). Association for Computational Linguistics, Online, 7740–7754. <https://doi.org/10.18653/v1/2020.emnlp-main.623>
- [22] M. Alexandra Kredlow, Michelle C Capozzoli, Bridget A Hearon, Amanda W Calkins, and Michael W Otto. 2015. The Effects of Physical Activity on Sleep: A Meta-Analytic Review. *Journal of Behavioral Medicine* 38 (2015), 427–449.
- [23] Philippa Lally, Cornelia HM Van Jaarsveld, Henry WW Potts, and Jane Wardle. 2010. How Are Habits Formed: Modelling Habit Formation in the Real World. *European Journal of Social Psychology* 40, 6 (2010), 998–1009.
- [24] Liliana Laranjo, Adam G Dunn, Huong Ly Tong, Ahmet Baki Kocaballi, Jessica Chen, Rabia Bashir, Didi Surian, Blanca Gallego, Farah Magrabi, Annie YS Lau, et al. 2018. Conversational Agents in Healthcare: A Systematic Review. *Journal of the American Medical Informatics Association* 25, 9 (2018), 1248–1258.
- [25] Jungeun Lee, Sungnam Kim, Minki Cheon, Hyojin Ju, JaeEun Lee, and Inseok Hwang. 2022. SleepGuru: Personalized Sleep Planning System for Real-life Actionability and Negotiability. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. ACM New York, NY, USA, 1–16.
- [26] Lihong Li, Wei Chu, John Langford, and Robert E Schapire. 2010. A Contextual-Bandit Approach to Personalized News Article Recommendation. In *Proceedings of the 19th International Conference on World Wide Web*. 661–670.
- [27] Zilu Liang and Bernd Ploderer. 2016. Sleep Tracking in the Real World: A Qualitative Study into Barriers for Improving Sleep. In *Proceedings of the 28th Australian Conference on Computer-Human Interaction*. 537–541.
- [28] Zilu Liang, Bernd Ploderer, Wanyu Liu, Yukiko Nagata, James Bailey, Lars Kulik, and Yuxuan Li. 2016. SleepExplorer: A Visualization Tool to Make Sense of Correlations Between Personal Sleep Data and Contextual Factors. *Personal and Ubiquitous Computing* 20 (2016), 985–1000.
- [29] Peng Liao, Kristjan Greenewald, Predrag Klasnja, and Susan Murphy. 2020. Personalized HeartSteps: A Reinforcement Learning Algorithm for Optimizing Physical Activity. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 1 (2020), 1–22.
- [30] Goran Medic, Micheline Wille, and Michiel EH Hemels. 2017. Short- and Long-Term Health Consequences of Sleep Disruption. *Nature and Science of Sleep* (2017), 151–161.
- [31] Mike A Merrill, Akshay Paruchuri, Naghme Rezaei, Geza Kovacs, Javier Perez, Yun Liu, Erik Schenck, Nova Hammerquist, Jake Sunshine, Shyam Tailor, et al. 2024. Transforming Wearable Data into Health Insights Using Large Language Model Agents. *arXiv preprint arXiv:2406.06464* (2024).
- [32] Susan Michie, Michelle Richardson, Marie Johnston, Charles Abraham, Jill Francis, Wendy Hardeman, Martin P Eccles, James Cane, and Caroline E Wood. 2013. The Behavior Change Technique Taxonomy (v1) of 93 Hierarchically Clustered Techniques: Building an International Consensus for the Reporting of Behavior Change Interventions. *Annals of Behavioral Medicine* 46, 1 (2013), 81–95.
- [33] Susan Michie, Maartje M Van Stralen, and Robert West. 2011. The Behaviour Change Wheel: A New Method for Characterising and Designing Behaviour Change Interventions. *Implementation Science* 6 (2011), 1–12.
- [34] Joanna C Moullin, Kelsey S Dickson, Nicole A Stadnick, Bianca Albers, Per Nilsen, Sarabeth Broder-Fingert, Barbara Mukasa, and Gregory A Aarons. 2020. Ten Recommendations for Using Implementation Frameworks in Research and Practice. *Implementation Science Communications* 1 (2020), 1–12.
- [35] Inbal Nahum-Shani, Shawna N Smith, Bonnie J Spring, Linda M Collins, Katie Witkiewitz, Ambuj Tewari, and Susan A Murphy. 2018. Just-in-Time Adaptive Interventions (JIT-Is) in Mobile Health: Key Components and Design Principles for Ongoing Health Behavior Support. *Annals of Behavioral Medicine* (2018), 1–17.
- [36] Maurice Ohayon, Emerson M Wickwire, Max Hirshkowitz, Steven M Albert, Alon Avidan, Frank J Daly, Yves Dauvilliers, Raffaele Ferri, Constance Fung, David Gozal, et al. 2017. National Sleep Foundation's Sleep Quality Recommendations: First Report. *Sleep Health* 3, 1 (2017), 6–19.
- [37] Jonathan M Peake, Graham Kerr, and John P Sullivan. 2018. A Critical Review of Consumer Wearables, Mobile Applications, and Equipment for Providing Biofeedback, Monitoring Stress, and Sleep in Physically Active Populations. *Frontiers in Physiology* 9 (2018), 743.
- [38] Nhat Pham, Tuan Dinh, Zohreh Raghebi, Taeho Kim, Nam Bui, Phuc Nguyen, Hoang Truong, Farnoush Banaei-Kashani, Ann Halbower, Thang Dinh, et al. 2020. WAKE: A Behind-the-Ear Wearable System for Microsleep Detection. In *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services*. 404–418.
- [39] Lukasz Piwek, David A Ellis, Sally Andrews, and Adam Joinson. 2016. The Rise of Consumer Health Wearables: Promises and Barriers. *PLoS Medicine* 13, 2 (2016), e1001953.
- [40] James O Prochaska and Wayne F Velicer. 1997. The Transtheoretical Model of Health Behavior Change. *American Journal of Health Promotion* 12, 1 (1997), 38–48.
- [41] Mashfiqui Rabbi, Min Hane Aung, Mi Zhang, and Tanzeem Choudhury. 2015. 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 (UbiComp '15)*. Association for Computing Machinery, New York, NY, USA, 707–718. <https://doi.org/10.1145/2750858.2805840>
- [42] James F Sallis, Neville Owen, E Fisher, et al. 2015. Ecological Models of Health Behavior. *Health Behavior: Theory, Research, and Practice* 5, 43–64 (2015).
- [43] Mourad Sarrouti, Asma Ben Abacha, Yassine Mrabet, and Dina Demner-Fushman. 2021. Evidence-based Fact-Checking of Health-related Claims. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (Eds.). Association for Computational Linguistics, Punta Cana, Dominican Republic, 3499–3512. <https://doi.org/10.18653/v1/2021.findings-emnlp.297>
- [44] Fred Shaffer and Jay P Ginsberg. 2017. An Overview of Heart Rate Variability Metrics and Norms. *Frontiers in Public Health* 5 (2017), 258.
- [45] Lucy Yardley, Bonnie J Spring, Helen Riper, Leanne G Morrison, David H Crane, Kristina Curtis, Gina C Merchant, Felix Naughton, and Ann Blandford. 2016. Understanding and Promoting Effective Engagement with Digital Behavior Change Interventions. *American Journal of Preventive Medicine* 51, 5 (2016), 833–842.