PERSONALIZING PUBLIC HEALTH: USING DEEP Q-LEARNING TO LOWER TYPE II DIABETES RISK IN INDIA

by

SABRINA REIS

A THESIS

Presented to the Department of Computer Science
and the Robert D. Clark Honors College
in partial fulfillment of the requirements for the degree of
Bachelor of Science

March 2023

Acknowledgements

This thesis is the product of many people's investment in me as a student, researcher, and person. I am deeply grateful to my thesis supervisor Thanh Nguyen for taking a chance on me by providing me with the opportunity to work on this project. You have provided an excellent example of the resilience, humor, and conviction needed to find fulfilment and success in research. I would also like to thank my mentors Michal Young and Weng-Keen Wong for their guidance as I navigated the writing and research process. Michal, you have always gone above and beyond to look after my happiness and well-being. I simply would not have made it this far without your support. Weng-Keen, I would not be the researcher I am today without your thoughtful mentorship. I must also thank Christabelle Dragoo and Denise Elder of the McNair Scholars Program for shaping my view of research as service. Finally, I thank the wonderful people managing the Clark Honors College Mentored Research Program and the UO VPRI Fellowship for funding for this work. I aspire to live up to the potential that all of these people see in me.

An Abstract of the Thesis of

Sabrina Reis for the degree of Bachelor of Science
in the Department of Computer Science to be taken March 2023

Title: Personalizing Public Health: Using Deep Q-Learning to Lower Type II Diabetes Risk in India

Approved: *Thanh H. Nguyen, PhD*
Primary Thesis Advisor

The high prevalence of type II diabetes in India necessitates novel public health tools to promote behaviors that lower the likelihood of developing diabetes and associated comorbidities. We introduce a public health text messaging system that learns from participant feedback to target each individual with messages that are most relevant to their current health behaviors. The participants who completed the AI-assisted text messaging intervention outperformed the control group in three out of eight categories. Highly engaged participants outperformed the control group in four out of eight categories and outperformed moderately engaged participants in seven out of eight categories, suggesting that engagement with the intervention is positively correlated with behavior change. We also perform a demographic analysis of the intervention results along the lines of sex, education level, and age. Recommendations for future work include personalizing message timing to increase intervention engagement and leveraging the plethora of data that is generated by the AI-assisted approach.

Table of Contents

List of Figures

1 Introduction

1.1 Research Motivation

One of the primary challenges to the advancement of public health goals worldwide is the prevalence of chronic, non-communicable diseases. This trend is particularly troubling given the preventable nature of these diseases in many people when provided with the resources needed to begin or maintain disease prevention behaviors like exercise. Non-communicable chronic diseases, in particular type II diabetes, are especially common in India. The harmful impact of type II diabetes on mortality and morbidity, or the development of additional health conditions, demonstrates the need for public health interventions that seek to reduce the prevalence and incidence of type II diabetes in India. We propose augmenting the traditional tactics used in public health interventions with artificial intelligence (AI) to maximize the effectiveness, accessibility, and flexibility of a type II diabetes intervention.

The motivation to integrate AI into public health interventions arises from its capacity to recognize patterns over time and make informed predictions. When applied to human behavior, these properties of AI may allow public health professionals to gain insights into the behavior patterns of individuals, enabling them to tailor a public health intervention to each person's needs. Given that public health interventions typically aim to reach the widest audience possible, public health outreach and messaging efforts must be sufficiently general so that they are relevant to a majority of their target demographics. Such generalization may prevent public health interventions from addressing individuals' values, concerns, or circumstances, which may limit the effectiveness of an intervention. Personalizing a public health intervention using AI may therefore improve the effectiveness of the intervention by increasing the relevance and, accordingly, the impact of public health messaging.

1.2 Research Questions

The potential impact of personalization on public health interventions naturally raises the question of whether personalized type II diabetes outreach is more effective at producing healthy behavior change than static outreach. The public health non-profit Arogya World investigated the static approach through a study in which all participants received the same set of text messages containing public health guidance on type II diabetes. These text messages encouraged higher levels of physical activity and consumption of fruits and vegetables. The study found that participants who received the texts "demonstrated greater improvement in a health behavior composite score" than participants who did not receive any messages and were less likely to exhibit a decline in health behaviors than the control group [18]. Among those who received the text messages, researchers observed "improved fruit, vegetable, and [saturated] fat consumption," but did not observe an increase in daily activity when compared with the control group [18].

Due to differences in study design discussed later in this work, it is not possible to make a direct comparison between the behavior change produced by the static intervention in the Arogya World study and our AI-assisted intervention. Given this constraint, we ask the more general question of whether personalized messages produce a significant improvement in health behaviors compared to a control group. We also explore the possible advantages of an AI-assisted intervention as opposed to a static intervention.

1.3 Hypothesis

In this project, we hope to observe that study participants who received personalized messages exhibit greater improvements in health behaviors than the individuals in the control group who did not receive any messages. If our prediction is correct, then an AI-assisted intervention may serve as a viable alternative to existing public health interventions that rely on static approaches. We also hope to find that the personalized messages are more effective at promoting physical activity, an area in which the static intervention was unsuccessful. Such an outcome may indicate that personalized outreach is more effective at promoting physical activity than static outreach.

1.4 Thesis Outline

The thesis begins with a background section that provides the public health context for the type II diabetes epidemic in India and elaborates upon the aforementioned static intervention carried out by Arogya World. Next, we present related work in deep learning that addresses the conceptual challenges of our message personalization efforts. We then proceed to the methodology section, in which we describe the framework of the AI-assisted public health intervention and the algorithm used in message selection. Afterward, we report the results of the AI-assisted intervention and analyze these results based on participant engagement levels and demographics. Finally, we discuss the implications of our results for the public health and AI research communities and recommend areas for further exploration.

2 Background

2.1 Type II Diabetes in India

India, like many other low- and middle-income countries, experiences a high prevalence of type II diabetes, a chronic health condition that interferes with the body's ability to regulate blood sugar [18]. In 2019, there were 77 million Indians with diabetes, the second-highest concentration of individuals diagnosed with diabetes in any country [19]. Around 90% of diabetes diagnoses are type II diabetes, leading epidemiologists to declare that India has a type II diabetes epidemic [19]. Given the association of type II diabetes with decreased quality of life and premature mortality, the prevalence of type II diabetes in the Indian population constitutes a public health crisis [19]. However, approximately 70% of India's population lives in rural areas, where access to health care is limited or nonexistent. Urban areas also exhibit low healthcare accessibility, as the high demand for care outstrips clinical resources in many cities [18]. Consequently, the Indian medical system does not currently have the resources to address the type II diabetes epidemic.

2.2 Diabetes Risk Management

A multitude of factors influence diabetes risk, including genetics, ethnicity, age, and abdominal adiposity, or the body's tendency to store fatty tissue around abdominal organs [19]. Fortunately, individuals can adopt protective health behaviors that lower their risk of diabetes or enable them to better manage their condition, resulting in increased quality of life [19]. Therefore, while not a substitute for access to medical care, public health interventions can lessen the care gap by assisting individuals in the adoption of protective behaviors.

2.3 Static Type II Diabetes Intervention

The importance of protective behaviors compelled the global health non-profit Arogya World to conduct a public health intervention in 2011 encouraging the adoption of protective behaviors in India [18]. Mobile phone uptake is high in India, with the National Family Health Survey for 2019-2021 estimating that 97% of urban households and 91% of rural households have access to mobile phones [16]. Accordingly, Arogya World carried out the diabetes intervention through text, developing a mobile phone outreach campaign called mDiabetes [18] [9]. Arogya World sourced their participants by collaborating with Nokia, a major phone company in India, to voluntarily enroll 1 million participants [18]. They then developed 56 text messages designed to "motivate improvement in diabetes risk behaviors and increase awareness about the causes and complications" and consulted with Indian consumers to make the messages culturally appropriate [18]. Their text messaging system sent messages statically—that is, in a pre-determined order—twice a week for 6 months [18].

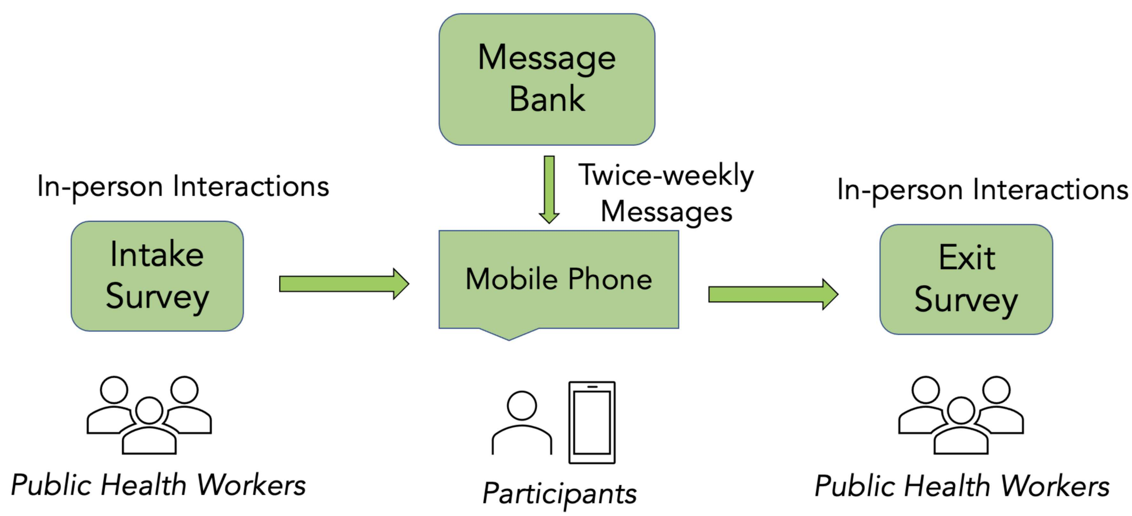

Figure 1: The structure of the static intervention conducted by Arogya World.

Researchers surveyed all participants twice, once prior to the text intervention and once more after the intervention had finished [18]. By comparing each participants' answers in the pre- and post-intervention surveys, the research team constructed a composite health score that considered exercise levels, moderation of fried food, consumption of at least 2 servings of fruit a day, and consumption of at least 2 servings of vegetables a day [18]. Each of these four categories received a score of 1, 0, or -1, where 1 indicated an increase in a health-promoting behavior over the course of the study, 0 represented no change, and -1 indicated a decrease in health-promoting behavior [18]. The scores for each category were then summed to create a composite health score that ranged from -4, signifying that a participant had decreased the frequency of all health-promoting behaviors, to +4, denoting that a participant had increased the frequency of all health-promoting behaviors [18].

The study reported that the test group and control groups had statistically significant differences in their composite health scores ($P < .001$), with test group participants more likely to receive higher composite scores and less likely to demonstrate a decrease in health-promoting behaviors over the study period compared to the control group [18]. Notably, while the test group exhibited greater improvements in fruit, vegetable, and fried food consumption, researchers did not find a significant difference in activity level improvement between the control group and test group [18]. These findings challenge the AI intervention to produce similar improvements in fruit, vegetable, and fried food consumption and promote higher levels of physical activity within the intervention group.

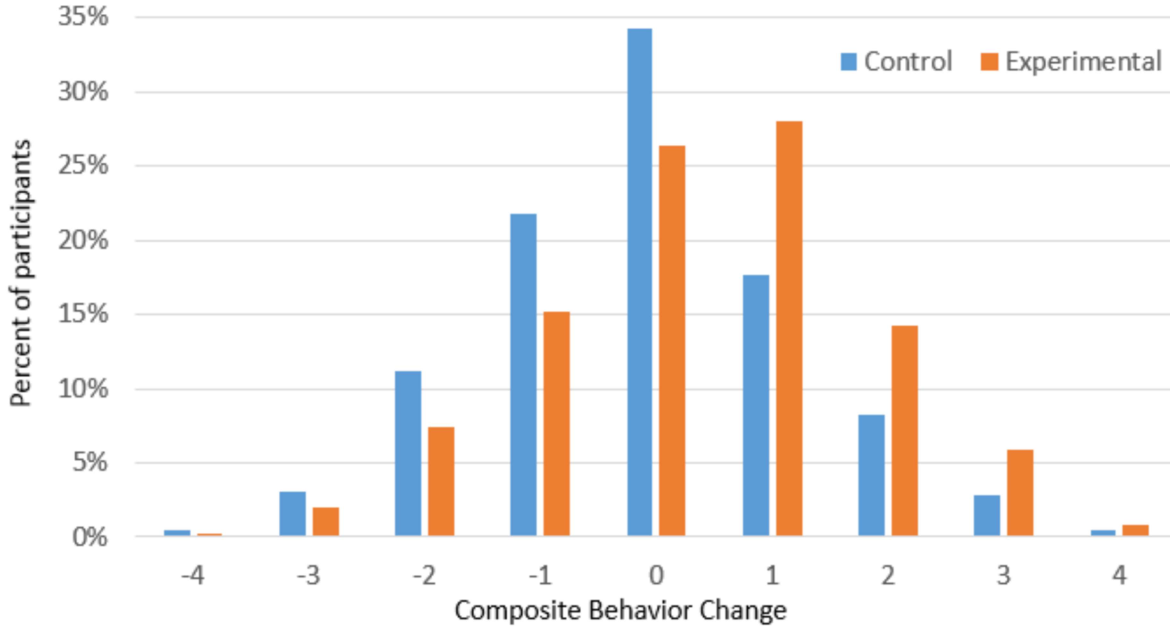

Figure 2: Composite health scores for the control and experimental groups. [18].

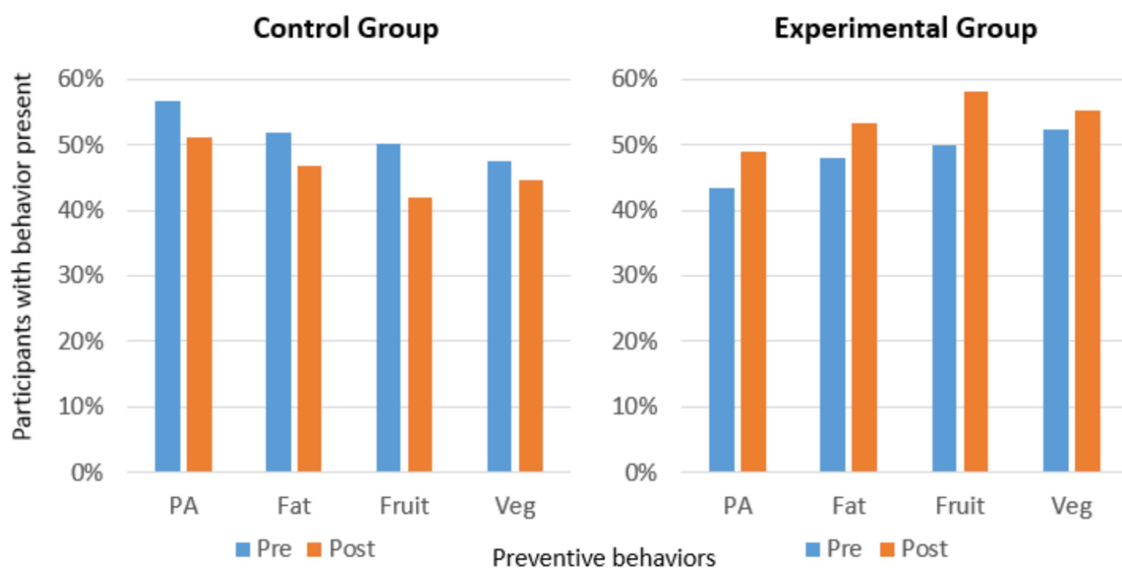

Figure 3: Percent of participants meeting behavior benchmarks in the control and experimental group. [18].

3 Related Work

3.1 Deep Q-learning for Decision Optimization

3.1.1 The Partial Observation Problem

The primary challenge facing the message selection model in the AI-assisted intervention is the lack of complete information on the state of a participant's health behaviors, as participants are not under constant observation. In other words, the true state of participants' health behaviors is only partially observable. Consequently, the message selection model must learn to make optimal message choices for each participant given incomplete information. Because real-world environments often involve dozens, if not hundreds, of variables that may affect the state, incomplete state information is common in applied machine learning (ML). The literature reviewed below establishes that a deep Q-learning network (DQN) is the state-of-the art technique for deep learning when states are only partially observable.

3.1.2 Deep Q-Learning for Partially Observable Processes

When a machine learning model does not have access to complete state information, the agent must gather information through trial-and-error decision-making. The observations that the agent makes in this process are fed into a learning policy which, over time, allows the agent to generate an estimation of its environment. Often, shallow learning methods, which only use one or two hidden layers in a neural network to detect patterns in data, are not complex enough to generate an accurate representation of the environment. Deep learning is typically more appropriate for these cases, as the many hidden layers that extract relevant information from the input data enable a model to "learn control policies in a range of different environments with only very minimal prior knowledge" [15]. DQNs are particularly well-suited to this task.

DQNs provide several distinct architectural advantages over regular Q-learning and other deep learning networks. First, DQNs employ a technique known as experience replay in which the model randomly selects a batch of previous timesteps that, along with data from the current timestep, are used to update the model [15]. Because DQNs require large amounts of data, the ability to reuse data reduces the difficulty of training, while random selection ensures that correlations between different timesteps are eliminated in the learning process [15]. Another key feature of DQNs is the use of a separate target network that serves as a copy of the DQN at a certain point in time [15]. The target network is periodically updated and used to train the main DQN, stabilizing the training process [15].

The advantages of these techniques are presented in "Human-level control through deep reinforcement learning,"which introduced the DQN architecture. Using a DQN, an agent learned to play 49 different Atari games without any prior knowledge, outperforming every other reinforcement learning method on 43 of the 49 games. The performance of the agent reveals the extent to which DQN can overcome deficiencies in state knowledge to make an accurate approximation of the true state. [15]. Additionally, the agent matched or surpassed the score achieved by a professional game tester on 29 of the games, which puts the inferential capacity of DQN on par with the human tester [15]. The games differed significantly in terms of the strategy needed to win, illustrating that DQN is a flexible approach suitable for many applications [15].

3.1.3 Limitations of Pure Deep Q-Learning

One limitation of pure DQN is that the algorithm only uses information from the last four states. If the algorithm needs to draw on information prior to the last four states, decisions made at the current timestep are no longer sufficiently independent from decisions made at earlier timesteps. In statistical terms, the process becomes non-Markovian, which means that a decision made by the agent at an earlier point in time may affect its decision at a much later timestep. This phenomenon violates a key assumption of DQNs, namely that the DQN is memoryless, meaning that the next state depends only on the current state [20]. To resolve the violation of the memoryless property, the authors of "Deep Recurrent Q-Learning for Partially Observable MDPs" replaced the first layer of the aforementioned Mnih et al.'s deep Q-network with a recurrent layer [8]. The authors then tested the deep recurrent Q-network on variants of Atari games where the objects of interests, such as an asteroid, disappeared for several frames to force the model to rely on earlier frames to beat the game, thereby violating the memoryless property [8]. The paper concluded that the deep recurrent Q-network slightly outperformed the pure deep Q-network in these scenarios [8]. Such an approach may be useful when the last four states are not a useful approximation of the current state or when the performance gain from resolving the violation of the memoryless property merits a more complicated approach than DQN.

3.2 Reinforcement Learning in Healthcare

Most applications of reinforcement learning in healthcare involve automating medical diagnoses or adapting treatment regimes to individual patients' needs [25]. Comparatively, there are few studies that apply reinforcement learning in the area of health management. Two studies have examined the effect of text messages encouraging physical activity among type II diabetes patients [24] [10]. When shallow reinforcement learning techniques were used to personalize text messages, participants in both studies increased their physical activity levels more than the control group, which received static messages. However, both of these studies were small with only 20-30 participants.

In contrast with existing work, our study is much larger, with over 500 participants completing the AI-based intervention. We also examine type II diabetes management from a more comprehensive perspective than prior work and the Arogya World study by considering physical activity, health knowledge and nutrition. Additionally, we expand beyond simple reminder messages to a message/question/response model in which participants may choose to respond to a follow-up question after receiving an informative message. Finally, the use of a DQN rather than a shallow approach enables the machine learning model to learn from participant data as efficiently as possible.

4 Methodology

4.1 Text Messaging Outreach

The AI-assisted public health intervention conducted by Arogya World surveyed approximately 1000 voluntarily enrolled participants over a six month period. The non-governmental organization Head Held High assisted in the selection of the 1000 participants by drawing upon their network of community leaders from villages across India who recruited participants, obtained informed consent, and collected their mobile phone number. Twice a week, participants received an informational message about diabetes prevention that the message selection model selected from a message bank. Afterward, the participants received a question message asking them to report changes in behavior. The optional responses to these questions are then used to update the message selection policy. There are five different types of messages: physical activity, cause knowledge, complication knowledge, healthy food intake, and risky food and substance intake.

4.2 Capturing Health Behaviors with an Ordinal Framework

As in the static intervention, Arogya World conducted an intake survey to establish a baseline for participants' health behaviors. The survey examined the five following categories of behaviors:

1. **Healthy food intake:** considers a participant's intake of foods included in a diet correlated with a lower risk of type II diabetes, including vegetables, fruits, whole grains, nuts, and legumes.

2. **Risky food and substance intake:** considers the frequency with which participants smoke, drink, and consume foods correlated with increased risk of

developing type II diabetes, such as fried food, fast food, tobacco products, and alcohol.

3. **Activity level:** includes all movement in a participant's day-to-day life, including housework, walking, sports, and formal exercise.

4. **Cause knowledge:** refers to a participant's awareness of the factors that are associated with an increased risk of type II diabetes. Knowledge encompasses both lifestyle factors and genetic factors, such as a family history of diabetes.

5. **Complication knowledge**: refers to a participant's awareness of how the development of type II diabetes affects a person's day-to-day life and how these complications can be managed.

Each question on the intake survey assessed one of the five categories of behaviors. Participants' responses to the questions were then classified according to how much protection their reported behaviors afforded against diabetes, with possible classes of 1 (low), 2 (medium), or 3 (high). We then aggregated the scores for questions examining the same category of behavior by computing the rounded average of their classifications, thereby classifying the overall level of protection afforded by a category of behaviors as 1 (low), 2 (medium), or 3 (high). To represent participants' health behaviors, we packaged these categorical classifications into a state vector for each participant. The entries in the state vector provide a snapshot of the state of participants' health behaviors. An example state vector is below.

$$\begin{pmatrix} 1 \\ 1 \\ 1 \\ 3 \\ 3 \end{pmatrix} : \begin{pmatrix} \textit{healthy food} \\ \textit{risky food and substances} \\ \textit{activity} \\ \textit{cause knowledge} \\ \textit{complication knowledge} \end{pmatrix}$$

From the state vector entries, we may infer that the participant's healthy food intake, risky food and substance intake, and activity level provide little protection against developing type II diabetes, though they are well-informed on diabetes causes and complications. While we may draw these narrow conclusions, it is important to remember that the true state of participants' health behaviors is partially observable, so the state vector can only provide an approximation of actual health behaviors. Consequently, conclusions about a participant's overall diabetes risk based on their state vector are limited.

4.3 Measuring Behavior Change

Responses to the twice-weekly questions were classified using the same scoring system of 1 (low), 2 (medium), 3 (high). Because each question addresses a single category of behavior, the weekly responses allow us to update the classification for a category of behaviors to reflect the participant's most recent habits. Accordingly, the entry in the feature vector that matches the category of behavior examined by the weekly question is updated to match the classification of the participant's response. In this way, we use changes in the attributes of participants' state vectors to capture real-world behavior changes.

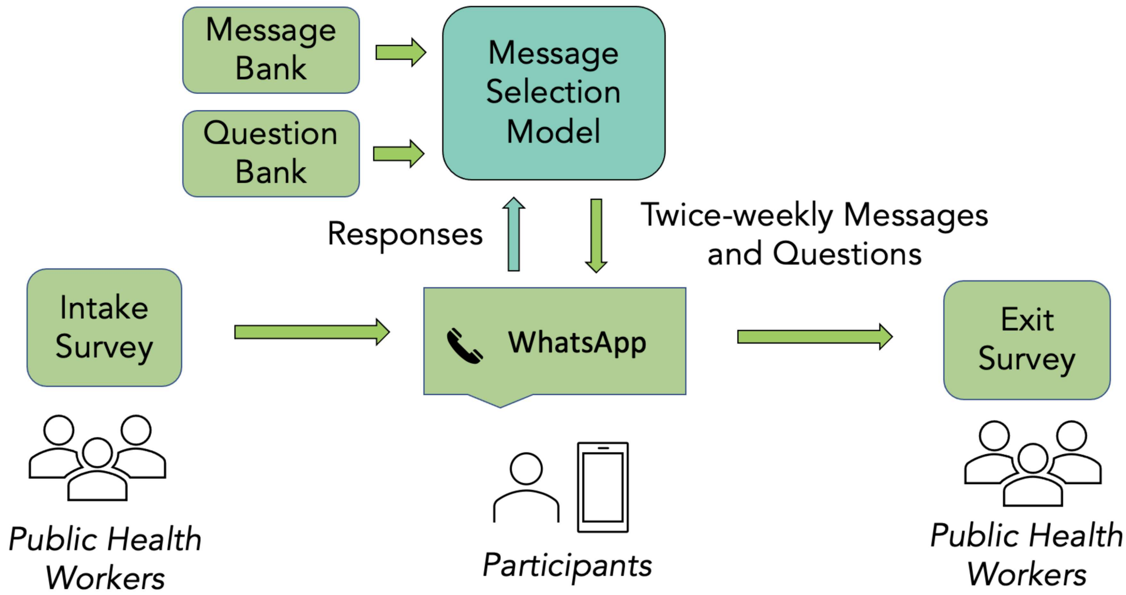

Figure 4: The structure of the AI-assisted intervention.

4.4 Message Selection Policy

To personalize message selection for each participant, we employ a DQN, a type of deep reinforcement learning model that focuses on the quality of different courses of actions in decision-making. In a machine learning context, quality refers to the ability to produce a desired outcome. In each state, or set of circumstances, a DQN must compare the quality of the actions available by computing the Q-value according to a policy, or pre-determined formula. The model will then choose the action with the highest Q-value, thereby optimizing decision-making outcomes in the long-term.

In this project, the different messages that may be sent to participants function as the actions that the DQN model must choose between, while the favorable outcome is a positive behavior change in the participant. Under the message selection policy, a message that appears more likely to induce a positive behavior change has a higher Q-value than other messages and will therefore get selected over less effective messages. The model is also rewarded for exploratory behavior so that, rather than selecting messages that address the same category of behavior each time, the model uses trial and error to identify as many potentially valuable messages as possible.

As mentioned in the Related Work section, a pure DQN model must contend with the fact that, in scenarios where the current state depends on states more than four timesteps prior, its performance may be inferior to that of a DQN equipped with a recurrent layer. Though a participant's current health behaviors are certainly related to their behavior many weeks prior, we determined that we can estimate the most relevant messages at the current timestep with sufficient accuracy using only the last four timesteps. Given this finding, the computational expense of adding a recurrent layer to the DQN was not justified, as any improvement would be marginal at best. Accordingly, our DQN has the standard DQN architecture.

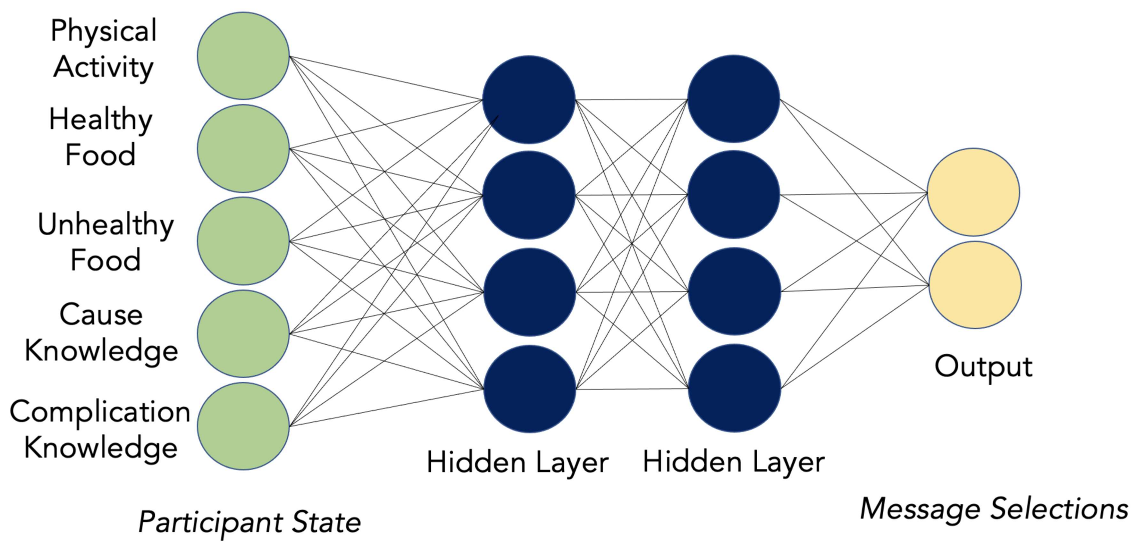

Figure 5: The Deep Q-Network used to select messages for participants.

5 Results

We first examine the efficacy of the AI-based intervention in provoking behavior change among all AI-assisted intervention participants. We then compare the impact of the AI-assisted intervention on participants of varying engagement levels. Finally, we perform a demographic analysis to determine how the results of the AI-based intervention differed on the basis of sex, education level, and age.

5.1 Comparison Between All AI Intervention Participants and Control Group

Of the 548 people that completed the initial intake survey for the AI-based intervention, 480 people finished, resulting in a dropout rate of approximately 12.4%. In comparison, 506 people completed the intake for the control group, and 441 of these people completed the outtake survey, leading to a dropout rate of approximately 12.8%. The similarity in the dropout rates between the intervention and control groups indicates that the intervention was well-tolerated.

In comparing all participants in the AI-based intervention to the control group, we found that AI-based intervention participants outperformed the control group in three out of eight subcategories: daily average exercise time, incidental exercise, and high fat food avoidance.

5.1.1 Composition of All AI Participants and Control Group

There are almost no observable differences in the makeup of the AI intervention group and the control group. The only exception is the age distribution of the participants; relative to group size, the control group contains more participants who are 18-25 and fewer participants who are 26-30 than the AI intervention group.

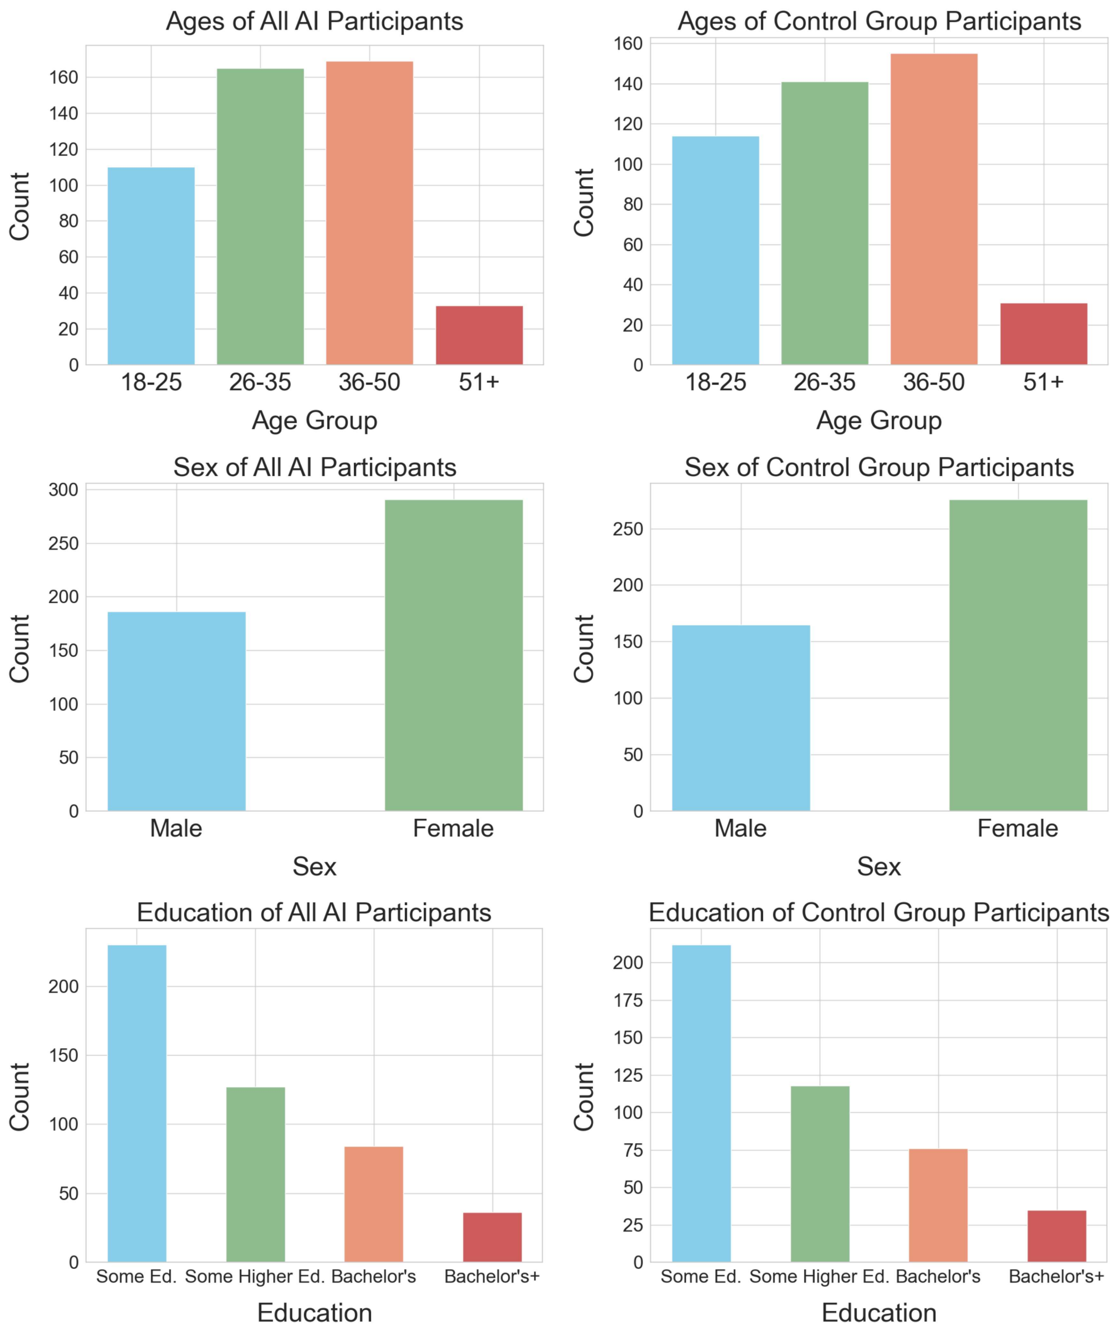

Figure 6: Demographics of all AI participants versus high responders.

5.1.2 Cause and Complication Knowledge

Participants in the AI group were more likely to improve their cause knowledge than participants in the control group. However, participants in the control group were slightly more likely to demonstrate an increase in complication knowledge.

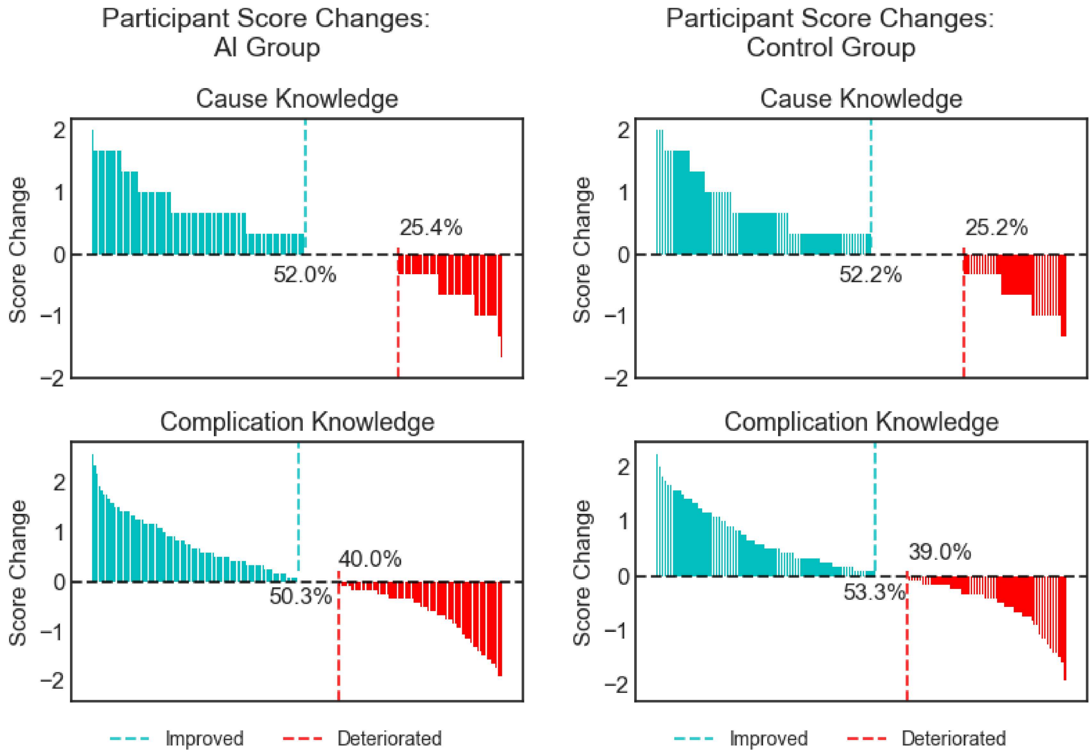

Figure 7: Knowledge score changes among all AI participants and control group participants.

5.1.3 Physical Activity

For analysis purposes, we separated physical activity into three categories: daily average exercise time, sports/workout/walking, and incidental exercise, which includes activity involved in day-to-day tasks such as farm work or climbing stairs. Participants in the AI group were more likely to increase their scores for daily average exercise time and incidental exercise than the control group. This result is especially notable given that the original static intervention did not observe a significant increase in physical activity among participants, suggesting that the AI-based model may be more effective at improving activity levels [18]. However, participants in the control group were more likely to improve their score in the sports/workout/walking category.

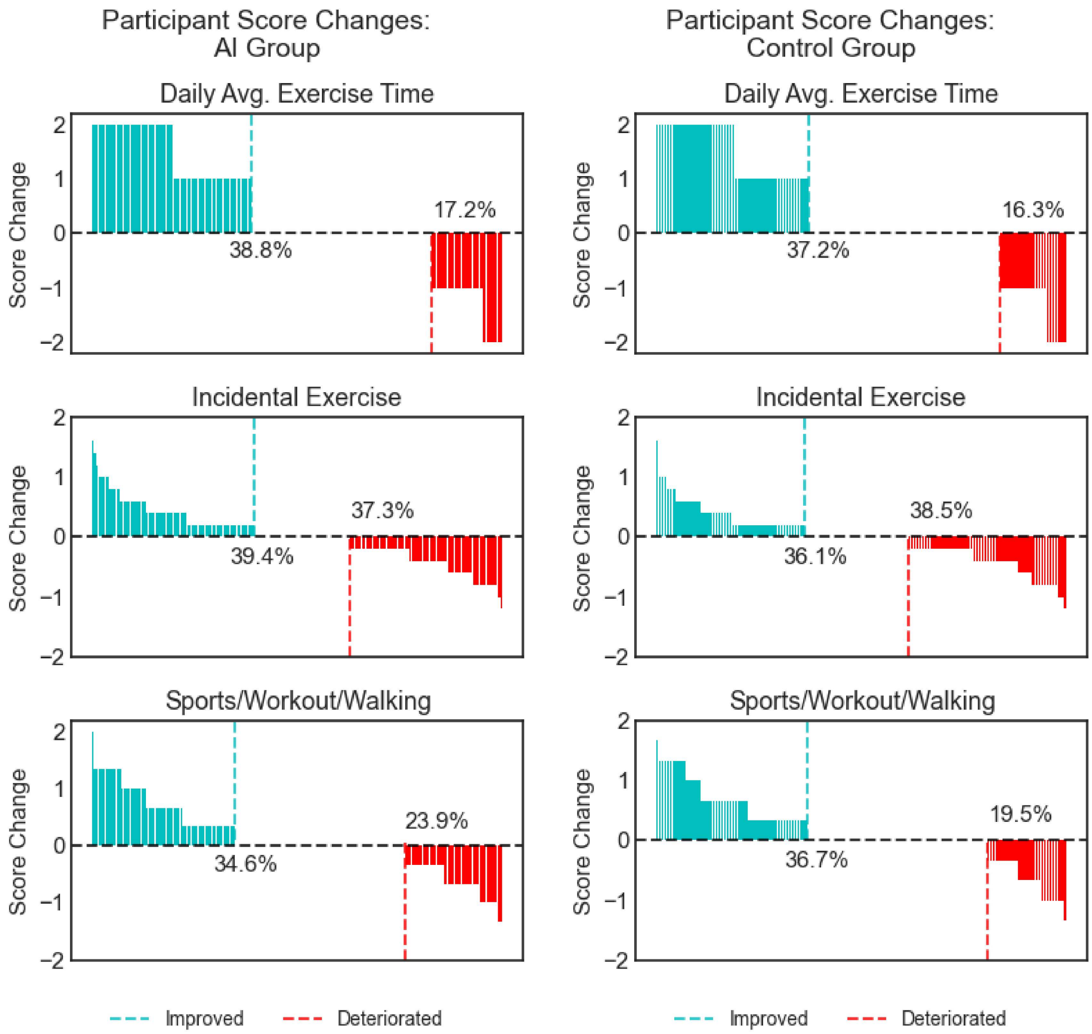

Figure 8: Physical activity score changes among all AI participants and control group participants.

5.1.4 Nutrition

As with the physical activity category, we divided the food intake category into high fat food avoidance, fruit consumption, and vegetable consumption to allow for detailed analysis. Participants in the AI group were more likely to impxrove their scores for high fat food avoidance, but less likely to increase their scores for fruit consumption and vegetable consumption compared to the control group.

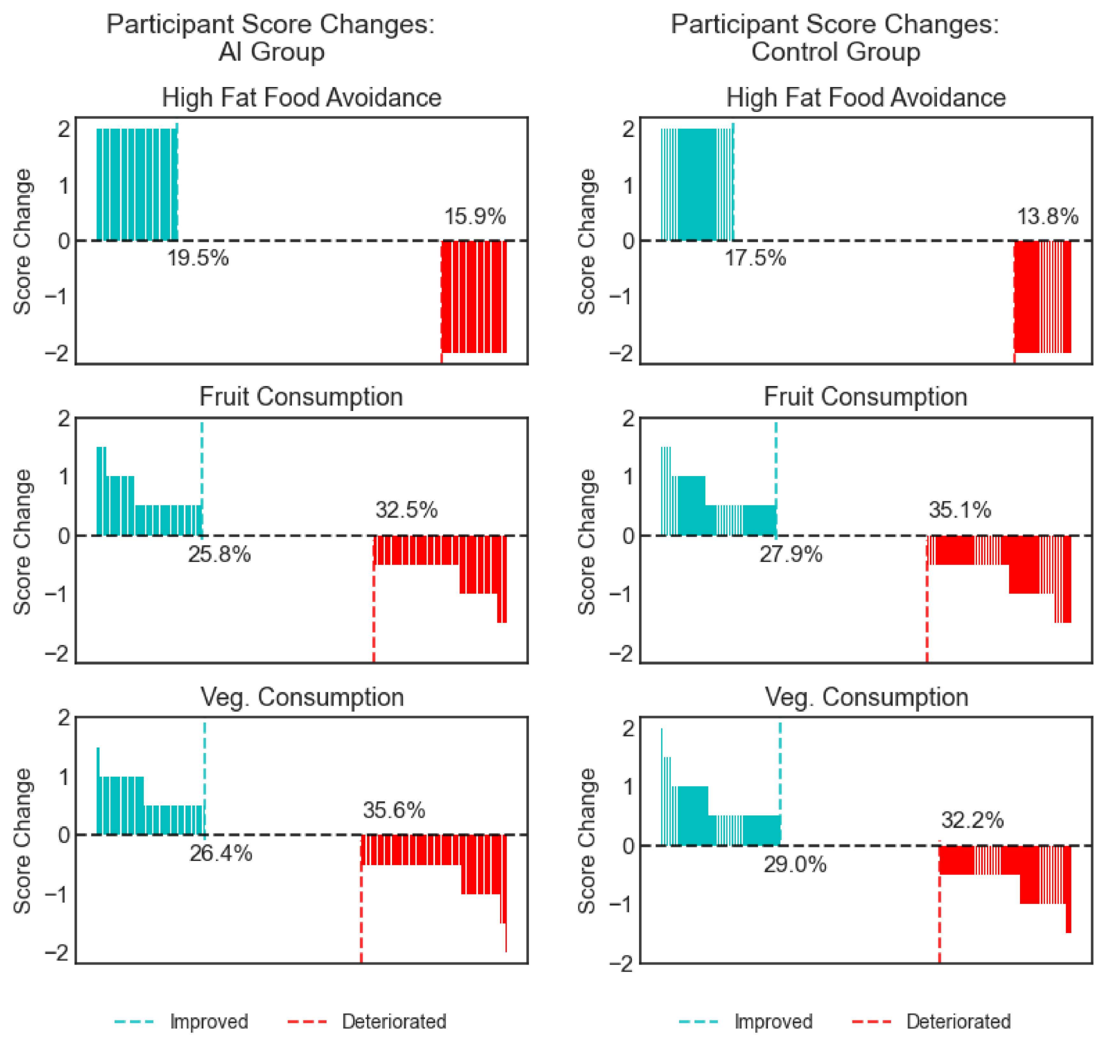

Figure 9: Nutrition score changes among all AI participants and control group participants.

5.2 Comparison Between High Responders and Control Group

To account for the low response rates of some participants, we compare the control group to participants in the AI-based intervention who answered at least 50% of the question messages. Because these participants were highly engaged with the intervention, they generated more data on their behavior changes over the course of the study, thereby providing more insight into the potential impact of the AI-assisted intervention compared to less engaged participants. Of the 480 participants who finished the AI-based intervention, 144 were high responders. These high responders outperformed the control group in four out of eight categories: cause knowledge, daily average exercise time, incidental exercise, and high fat food avoidance.

5.2.1 Composition of High Responders and Control Group

Demographically, the high responder group appears similar to the control group with a few notable exceptions. Proportionally, there are significantly more male participants in the high responder group than the control group. Additionally, participants in the high responder group are more likely to have completed at least some higher education compared to the control group.

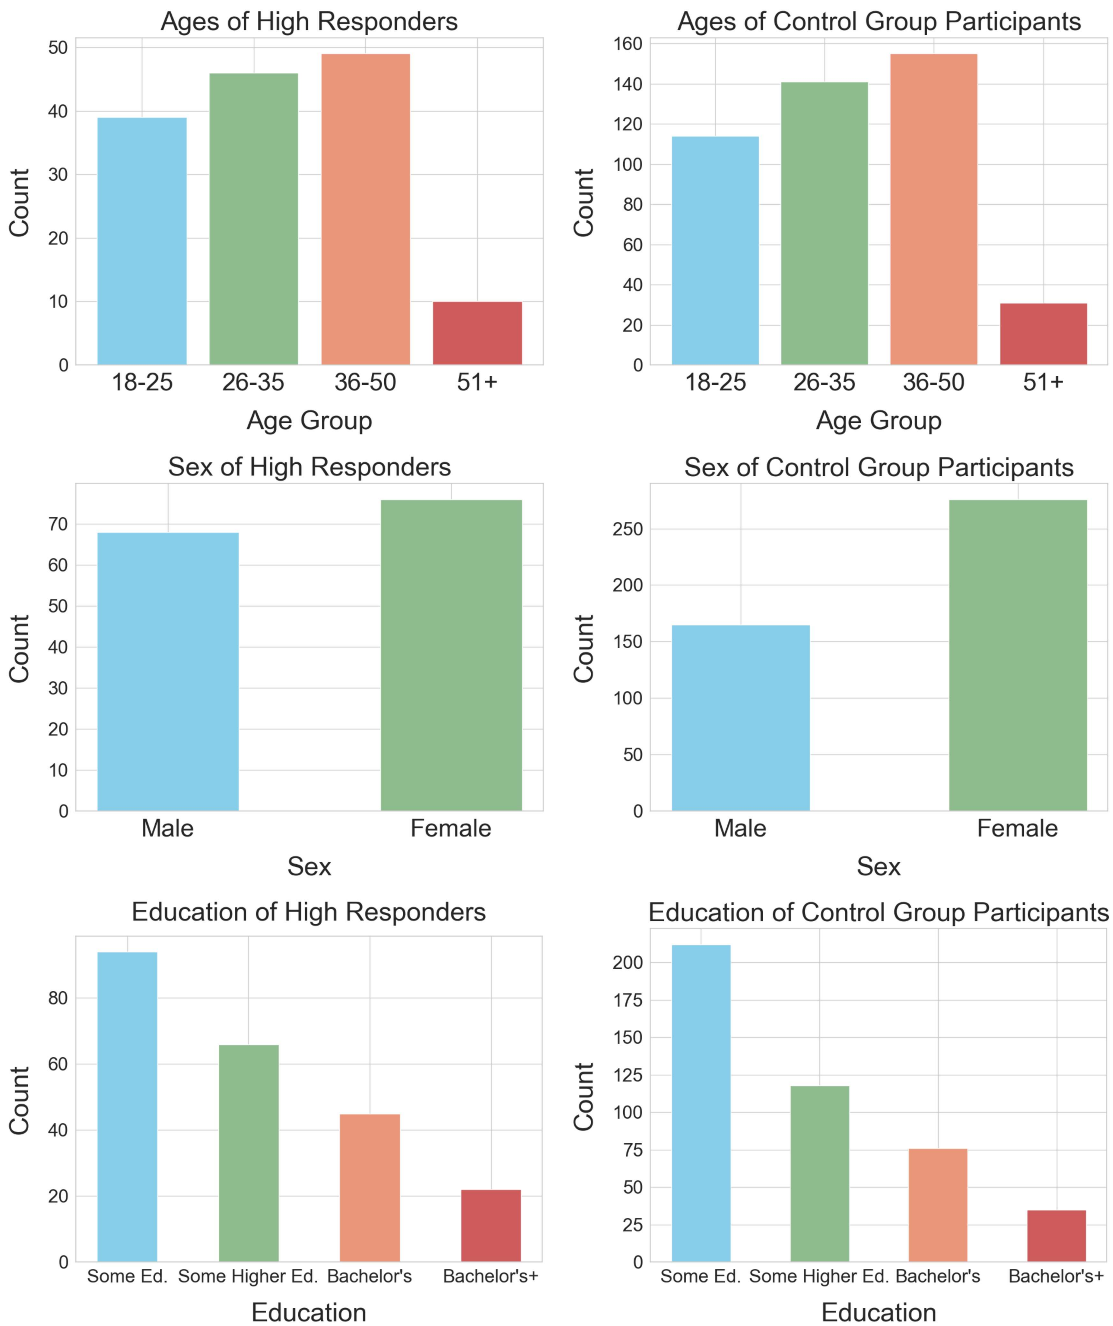

Figure 10: Demographics of the group that responded at least 50% of the time versus the control group.

5.2.2 Cause and Complication Knowledge

Participants in the AI-assisted intervention were more likely to improve their cause knowledge, while participants in the static group were more likely to experience an improvement in their complication knowledge. However, fewer participants in the AI-based intervention experienced a decline in their complication knowledge.

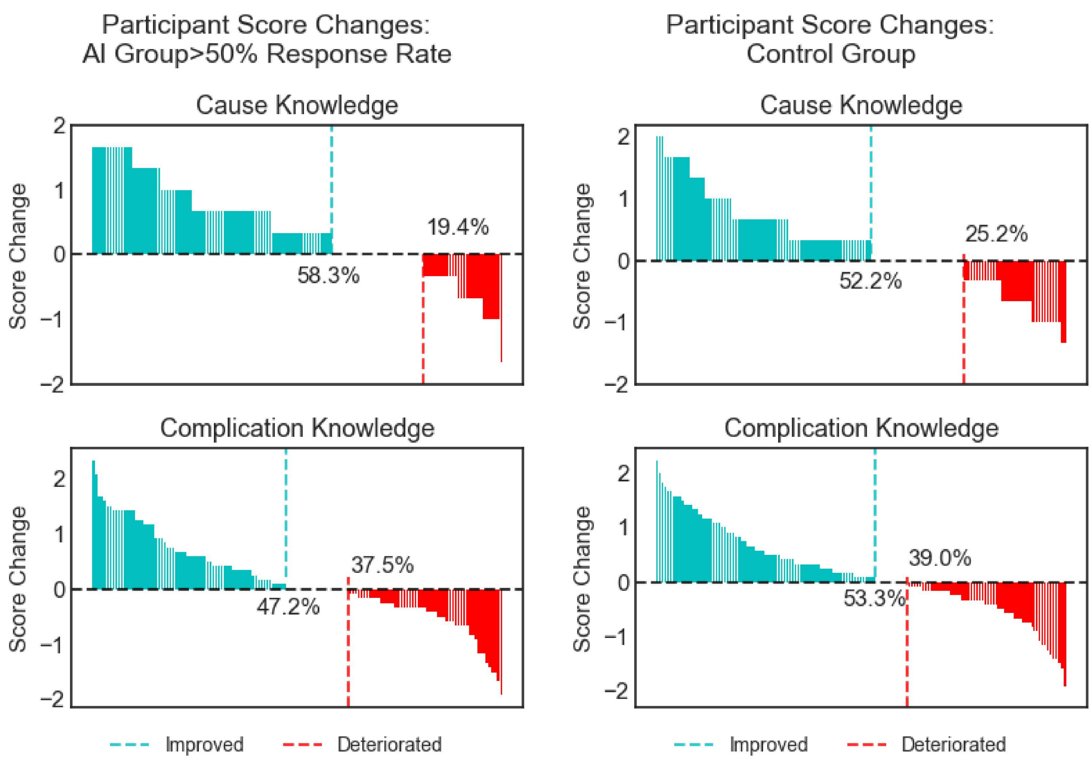

Figure 11: Knowledge score change among AI-based intervention participants and control group participants.

5.2.3 Physical Activity

Scores for daily average exercise time and incidental exercise were more likely to increase and less likely to decline in the AI intervention group. Participants in the control group were more likely to increase their scores in the sports/workout/walking category, but also more likely to experience a decline in their score for this category.

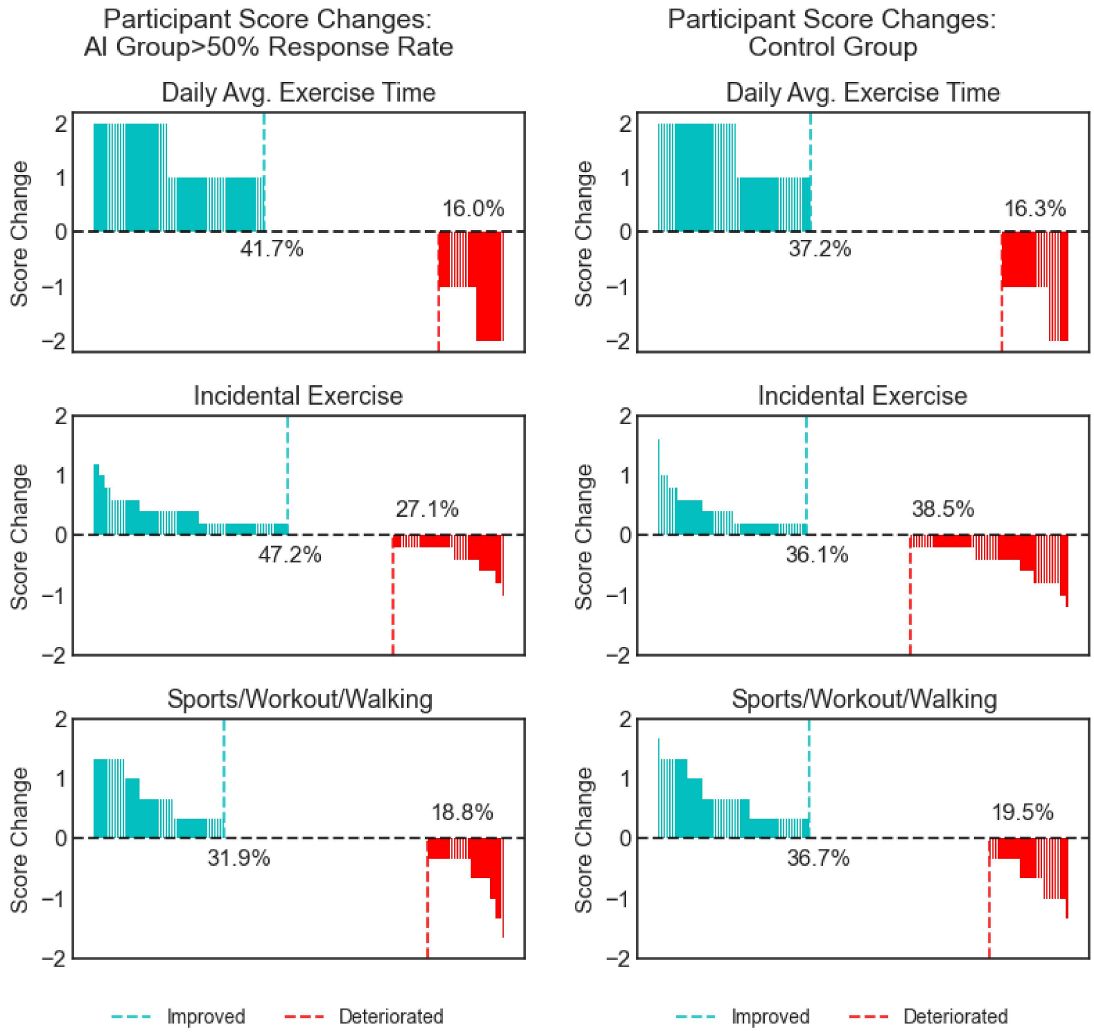

Figure 12: Physical activity score change among AI-based intervention participants and control group participants.

5.2.4 Nutrition

Compared to the static intervention group, the AI intervention participants exhibited a greater improvement in high fat food avoidance. In terms of fruit consumption, the static and AI-assisted intervention groups demonstrated comparable percentages of improvement, but the AI-assisted group experienced considerably less deterioration in their fruit consumption. Finally, with regard to vegetable consumption, the static group demonstrated slightly higher percentages of improvement and markedly lower percentages of deterioration compared to the AI group.

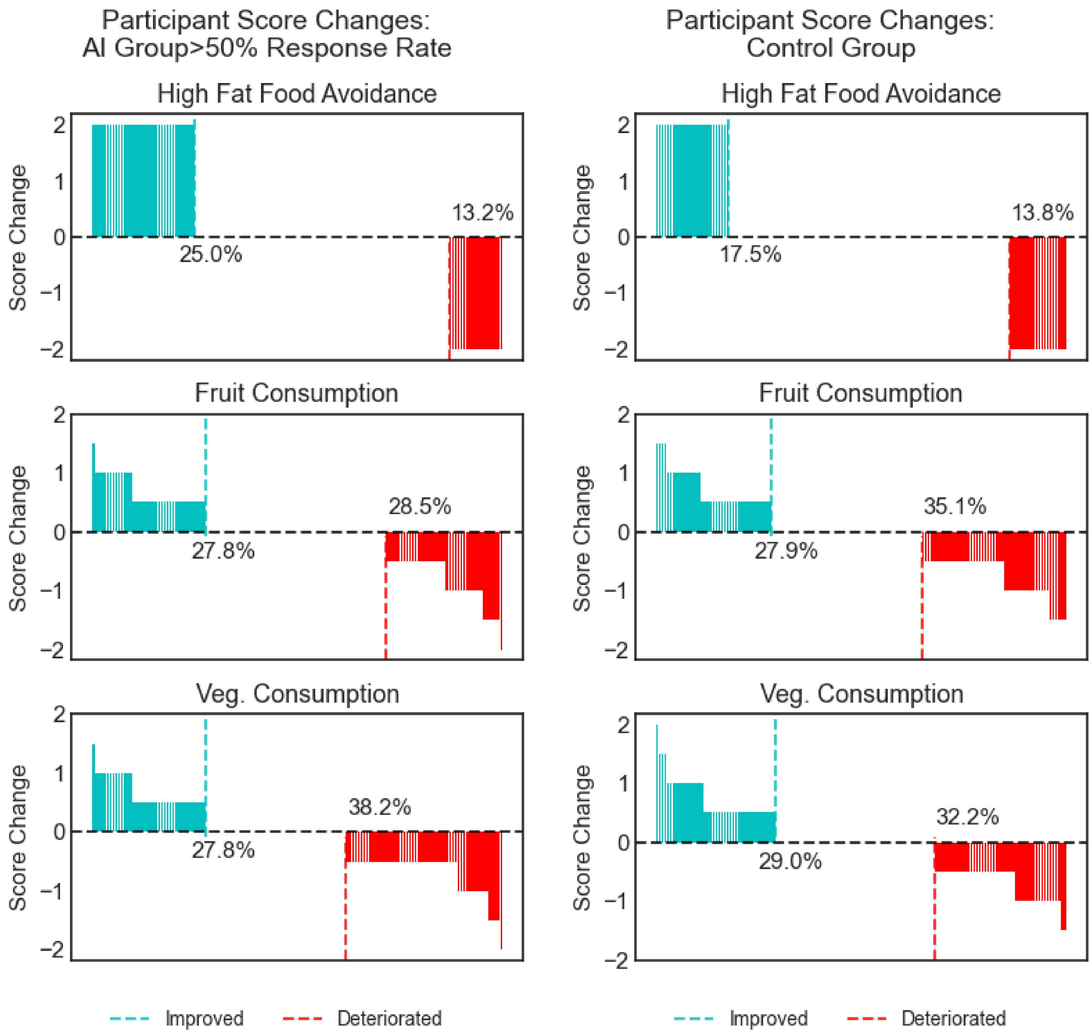

Figure 13: Nutrition score change among AI-based intervention participants and control group participants.

5.3 Comparison Between Moderate Responders and High Responders

We now compare the behavior changes among our high responders to the behavior changes among our moderate responders, or participants who had a response rate of at least 25%, to assess how the level of engagement with the intervention affects the degree of behavior change. The high responders outperformed the moderate responders in all categories except complication knowledge, suggesting that higher engagement is is correlated with a higher likelihood of improved behavior. This correlation provides further evidence that the AI-based intervention is effective at promoting behavior change.

5.3.1 Composition of Moderate Responders

Demographically, there are few observed differences between high responders and moderate responders. Compared to high responders, moderate responders are slightly more likely to be 26-30 and female.

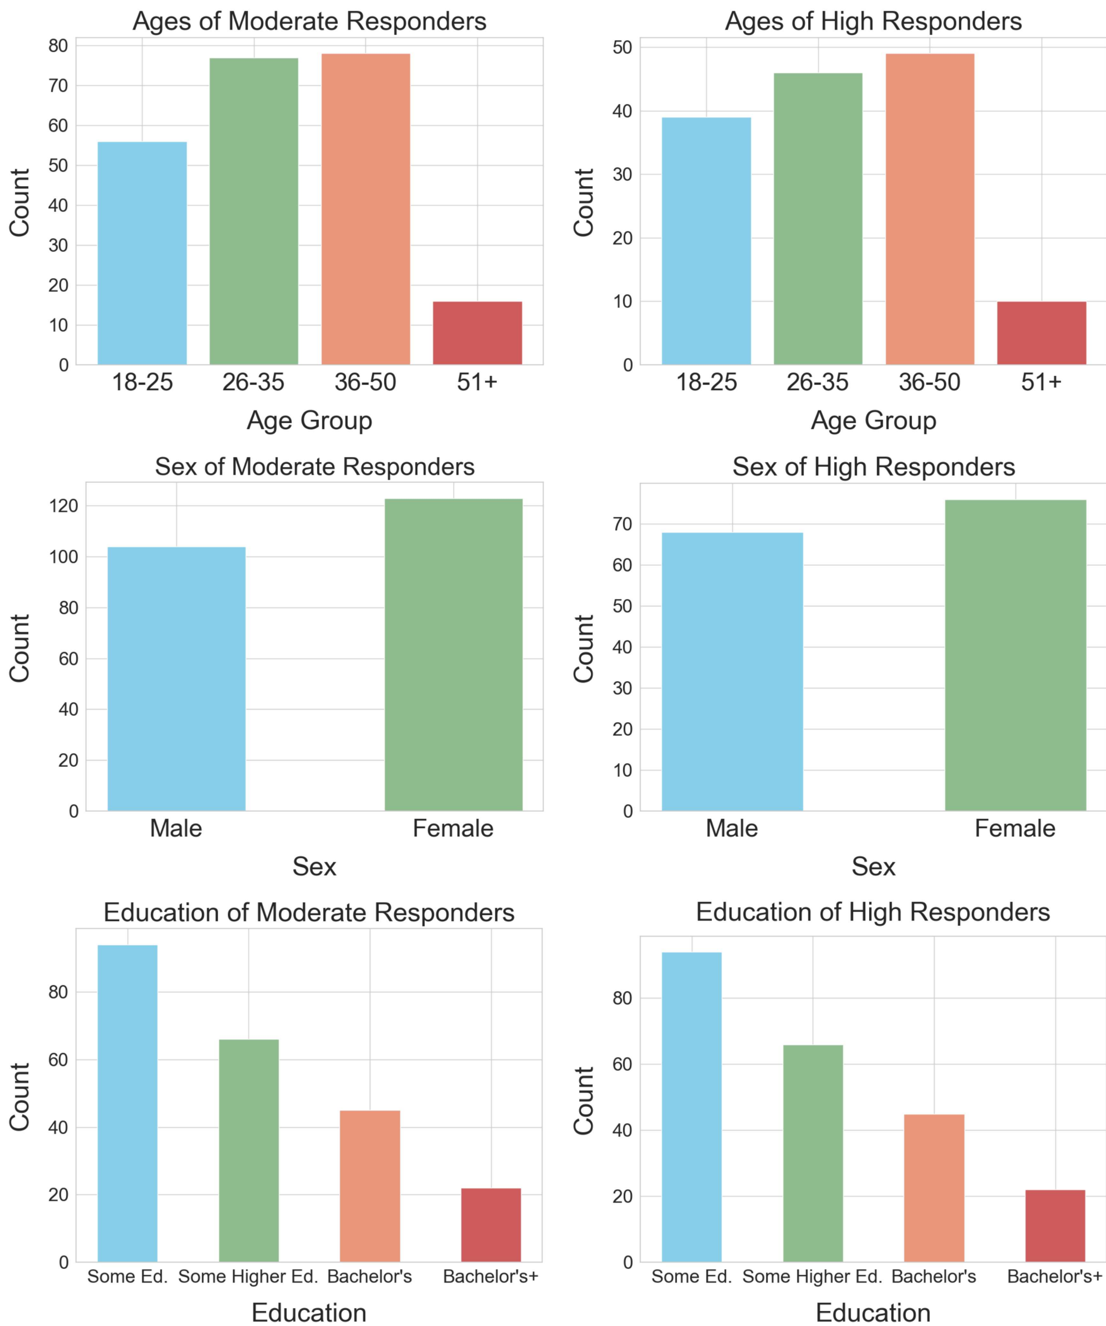

Figure 14: A demographic comparison of moderate responders versus high responders.

5.3.2 Cause and Complication Knowledge

High responders were more likely to improve and less likely to experience a decline in their cause knowledge score. However, moderate responders were marginally more likely to improve their complication knowledge score compared to high responders. This finding may indicate that increased engagement with the intervention does not lead to increased complication knowledge, which warrants further investigation.

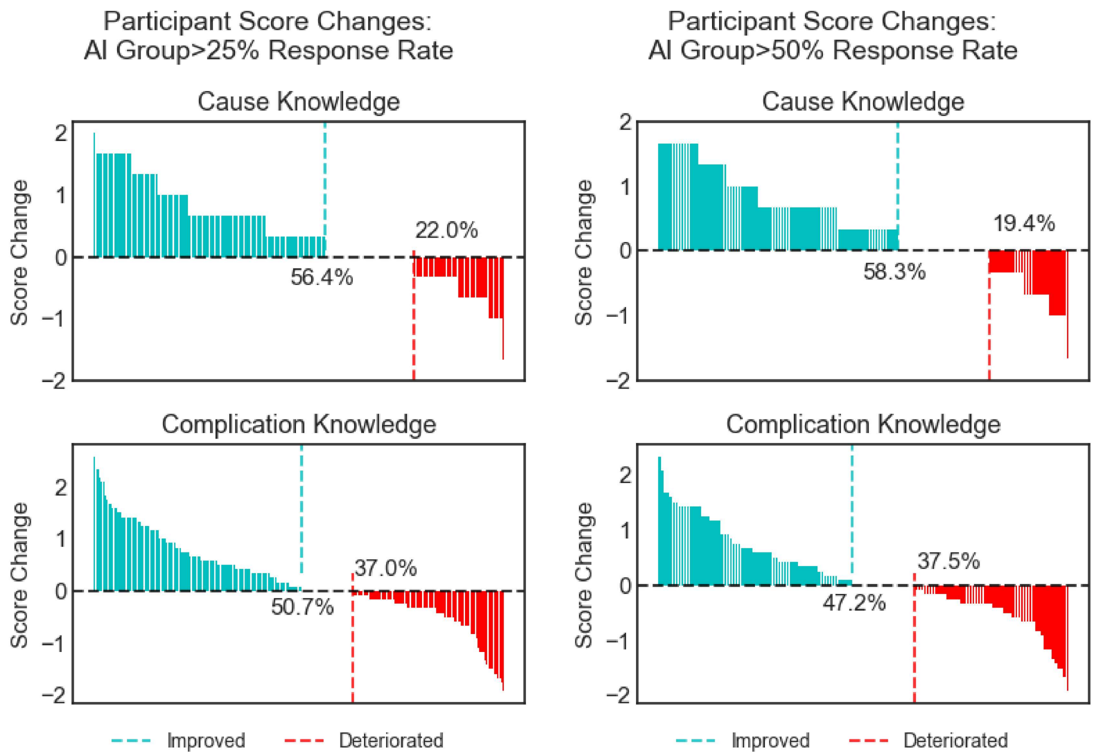

Figure 15: Knowledge score changes among moderate and high responders.

5.3.3 Physical Activity

Across all subcategories of physical activity, high responders were more likely to improve their scores and less likely to experience a decline. This result affirms that the AI-based intervention is effective at improving physical activity levels among engaged participants.

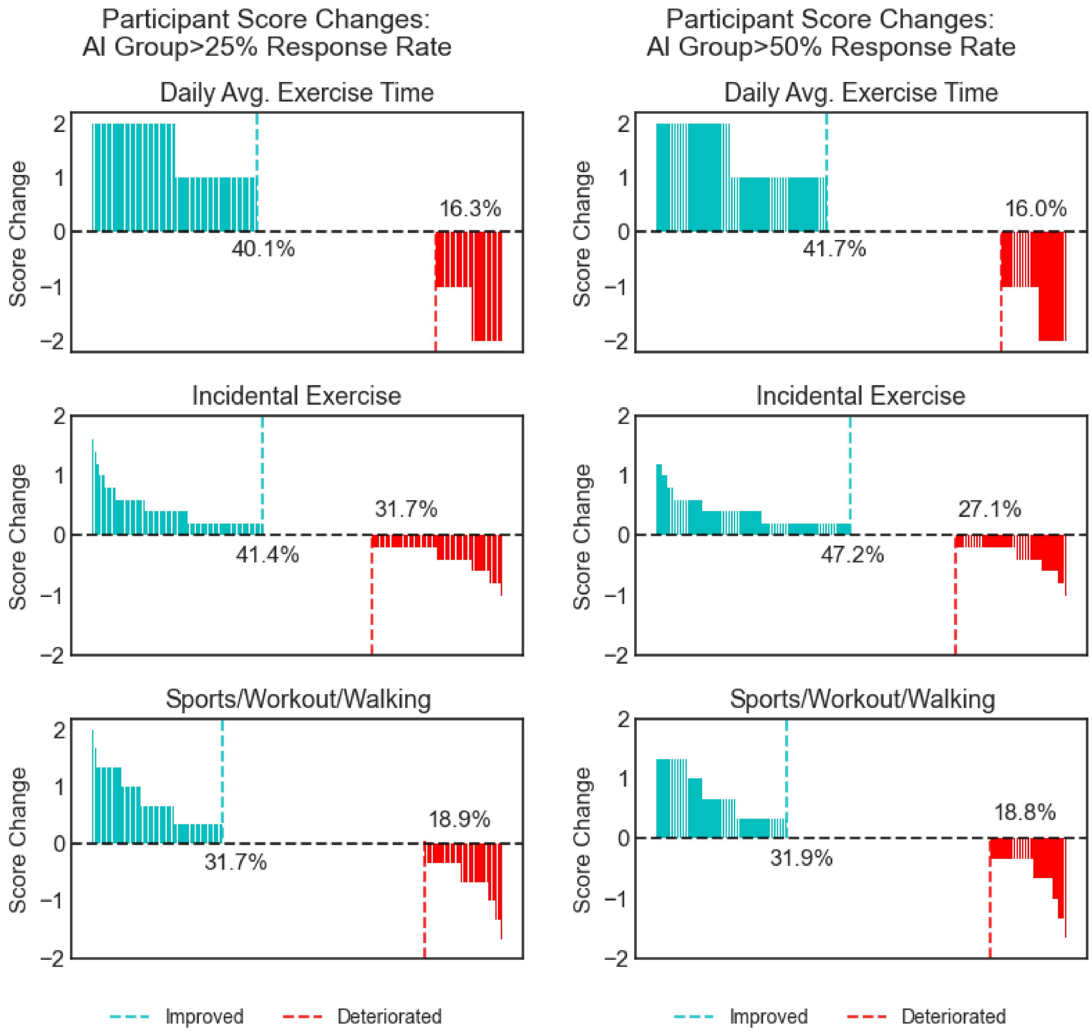

Figure 16: Physical activity score changes among moderate and high responders.

5.3.4 Nutrition

In all nutrition-related subcategories, high responders were more likely to improve their behavior scores and less likely to experience a decline compared to moderate responders. Once again, this finding demonstrates the effectiveness of the AI-based intervention in promoting behavior change among engaged participants.

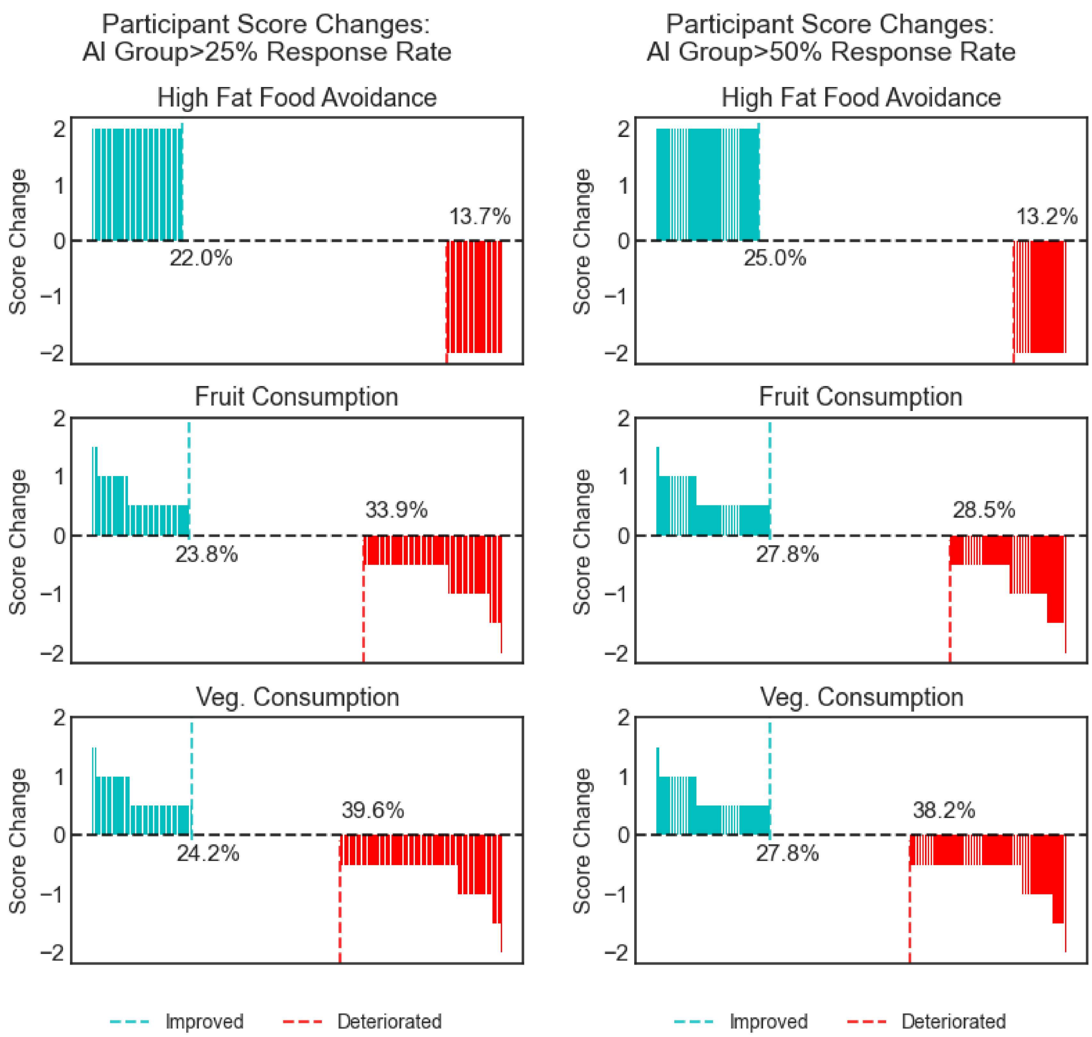

Figure 17: Nutrition score changes among moderate and high responders.

5.4 Demographic Analysis of Behavior Change

In addition to our previous analyses, we performed a demographic analysis of the AI-assisted intervention impact along the lines of sex, education level, and age. Different communities may struggle with different health barriers, and the demographic analyses provide valuable information on the behaviors of various groups for public health professionals. For these analyses, we compare participants in our high responder group, as this group provided the most data on their health behaviors over time.

5.4.1 Sex

Overall, male participants were more likely to improve their cause and complication knowledge scores and nutrition scores, while female participants were more likely to improve their physical activity scores.

Knowledge Female participants exhibited lower improvements in both cause and complication knowledge and higher rates of deterioration compared to male participants.

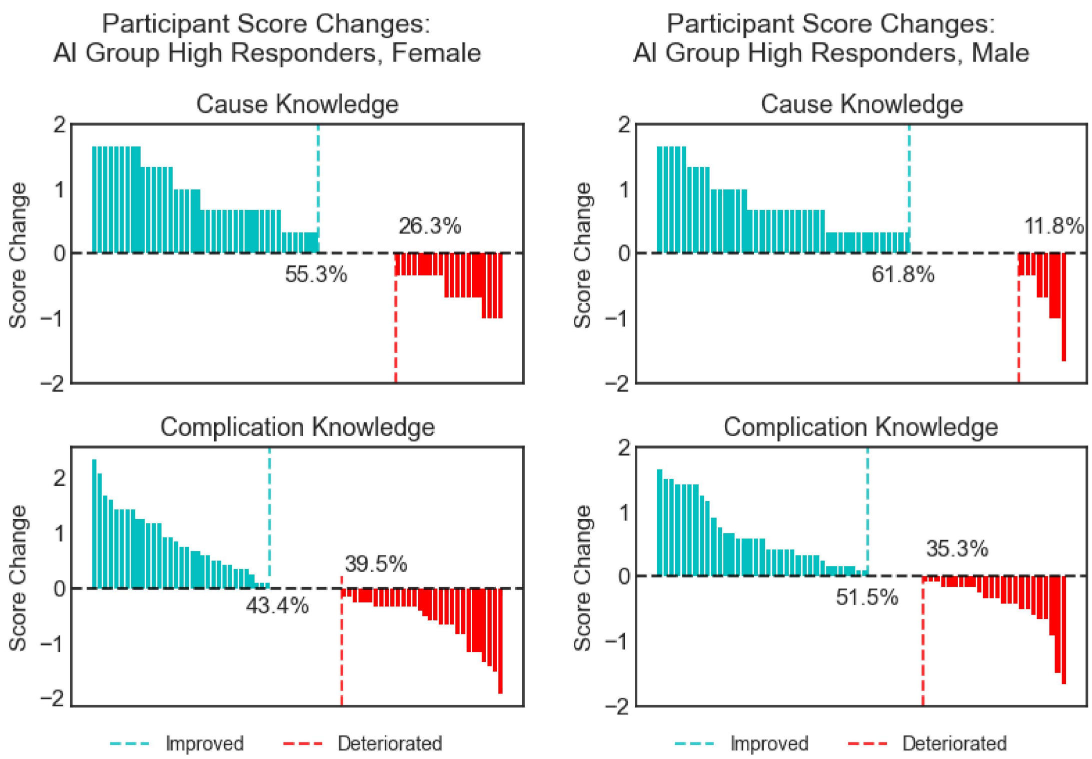

Figure 18: Knowledge score changes among female and male AI-based intervention participants.

Physical Activity Female participants experienced higher rates of improvement and lower rates of deterioration in terms of overall average exercise time. Much of this improvement in the female group comes from the combined sports/workout/walking category, which saw female participants improve much more than male participants. With respect to incidental exercise, female participants improved at rates similar to the male participants, but experienced much higher deterioration.

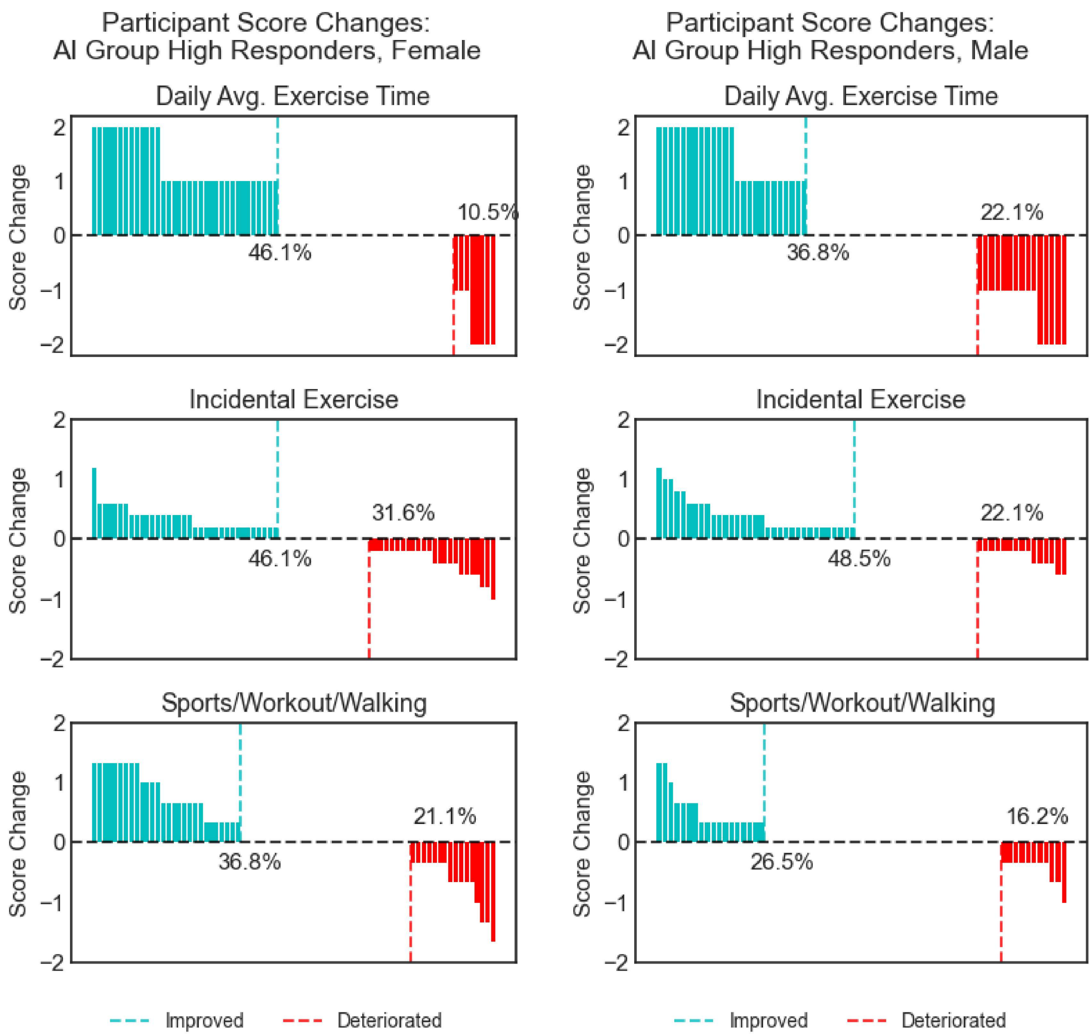

Figure 19: Physical activity score changes among female and male AI-based intervention participants.

Nutrition Across all subcategories, male participants were more likely to improve their nutrition-related scores. However, male participants were more likely to deteriorate in their fruit consumption compared to female participants.

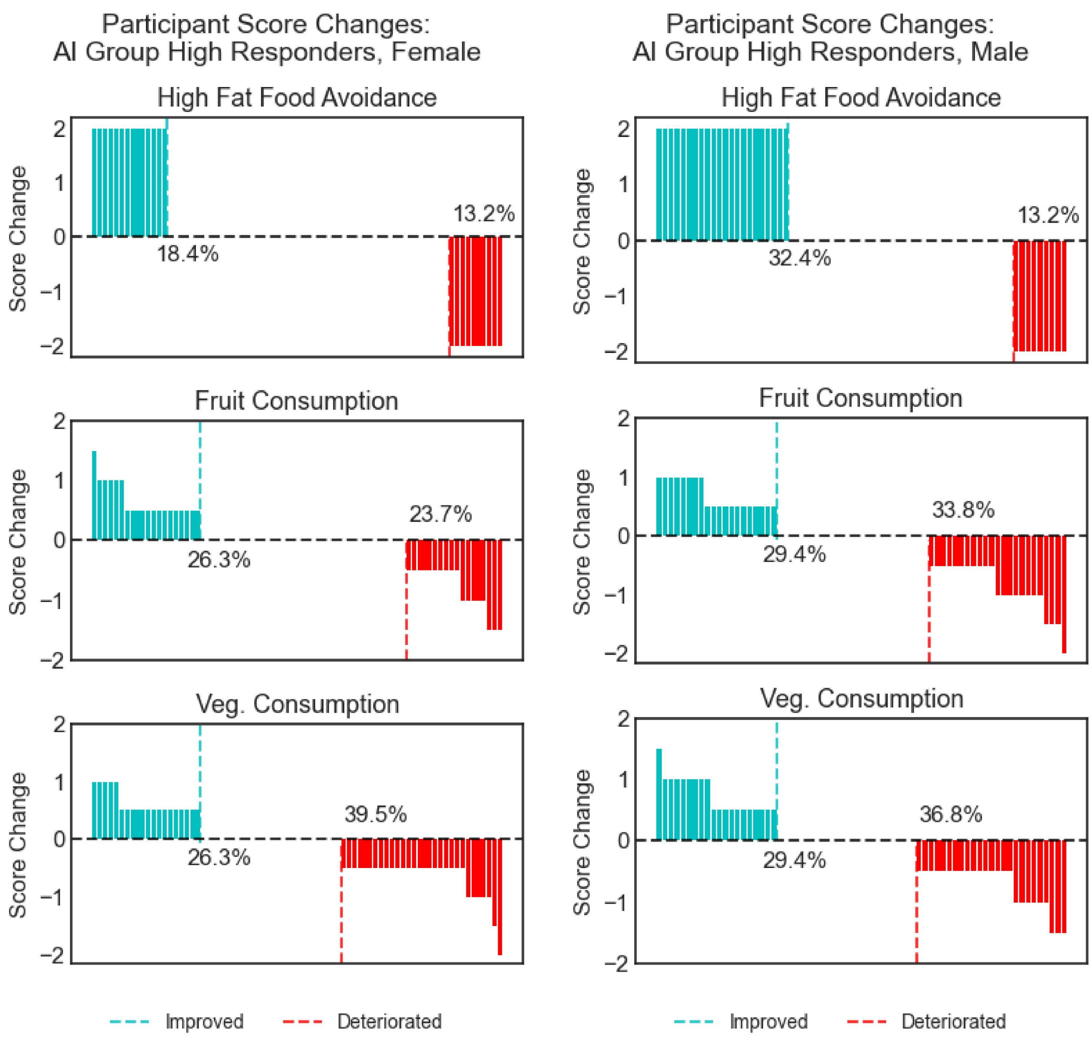

Figure 20: Nutrition score changes among female and male AI-based intervention participants.

5.4.2 Education

To analyze the relationship between education level and the effect of the intervention, we split participants up into two groups: those who had no higher education and those who had at least some higher education education, regardless of whether they earned a degree. Generally, participants with some higher education were more likely to improve their cause and complication knowledge and high fat food avoidance, while participants who had no higher education were more likely to improve their physical activity scores and fruit consumption.

Knowledge In both cause and complication knowledge, participants with no higher education were less likely to improve their scores and more likely to deteriorate. These results indicate that more work is needed to ensure that the knowledge messages are accessible to people who do not come from highly educated backgrounds.

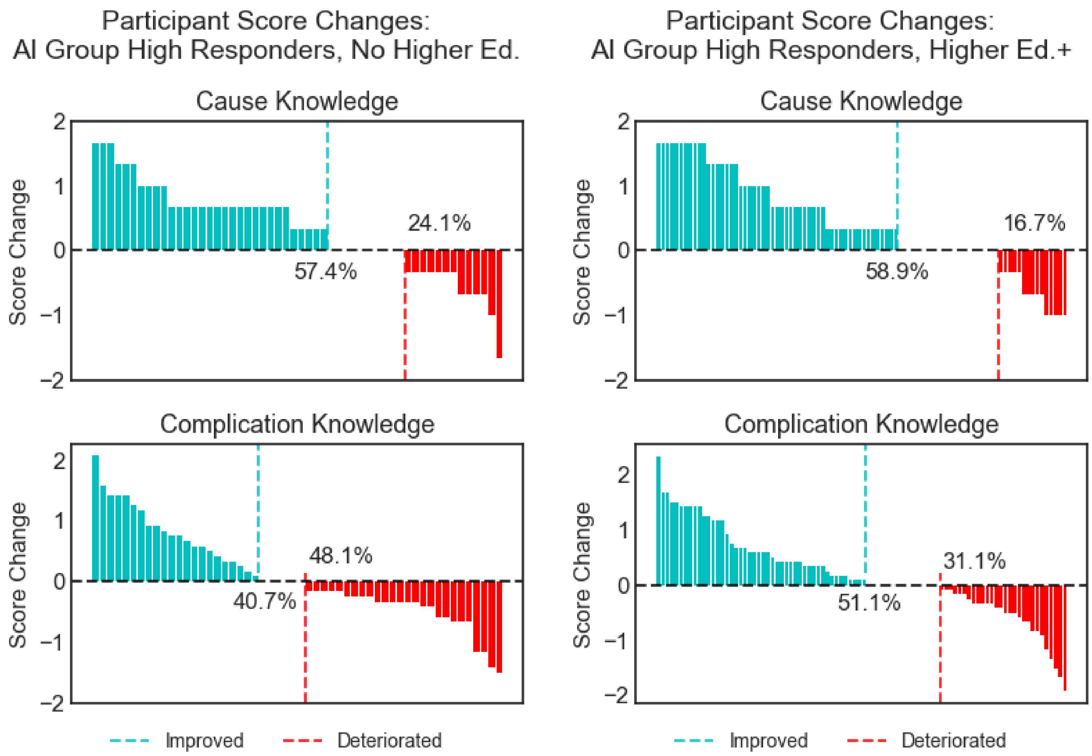

Figure 21: Knowledge score changes among AI intervention participants with and without some college education.

Physical Activity Across all subcategories, the group without any higher education was more likely to improve their physical activity-related scores and less likely to decline.

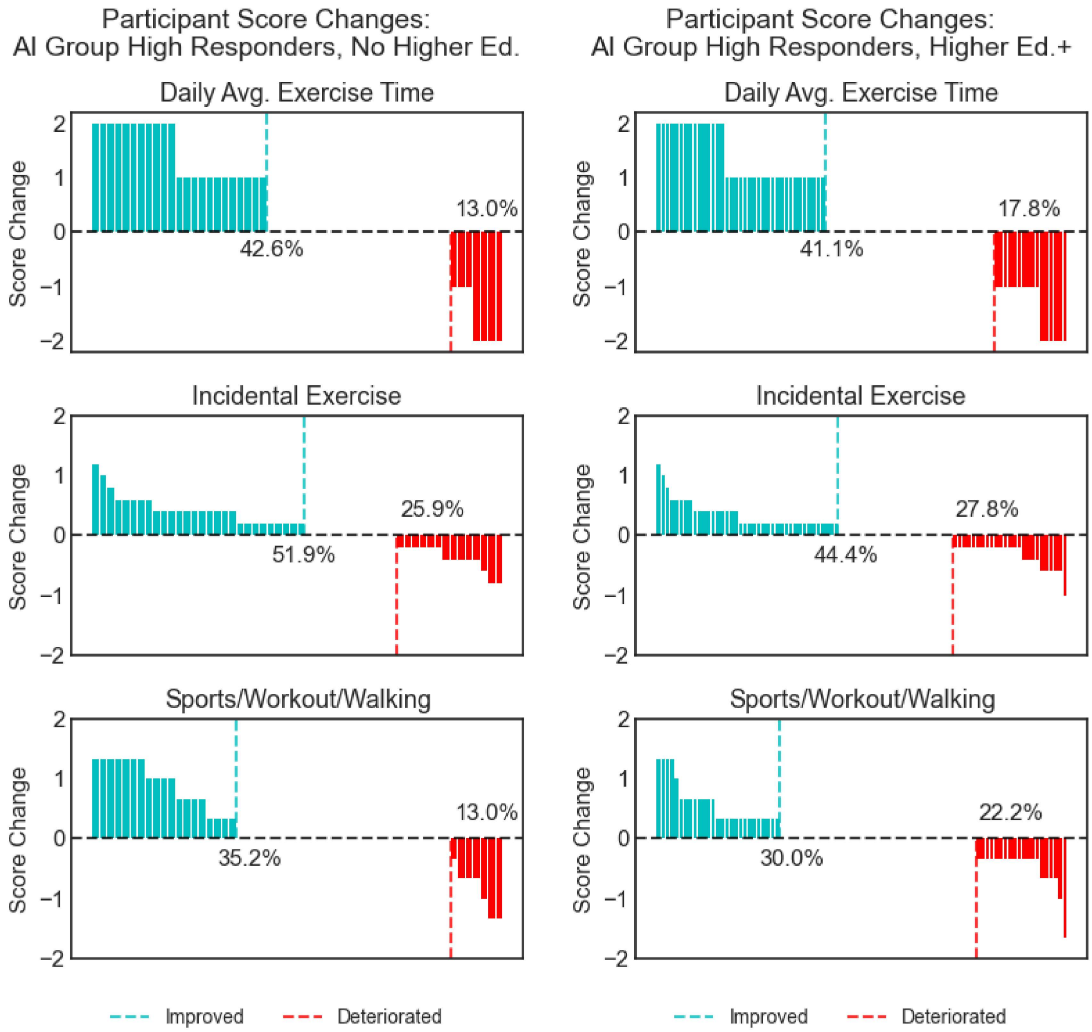

Figure 22: Physical activity score changes among AI intervention participants with and without some college education.

Nutrition In the high fat food avoidance category, participants with some higher education displayed markedly higher rates of improvement compared to the group without higher education. The higher education group was also more likely to display improvements in fruit and vegetable consumption. Rates of deterioration in the fruit and vegetable categories were mixed, however; fruit consumption deteriorated more among the group with some higher education, while vegetable consumption deteriorated more among participants without higher education.

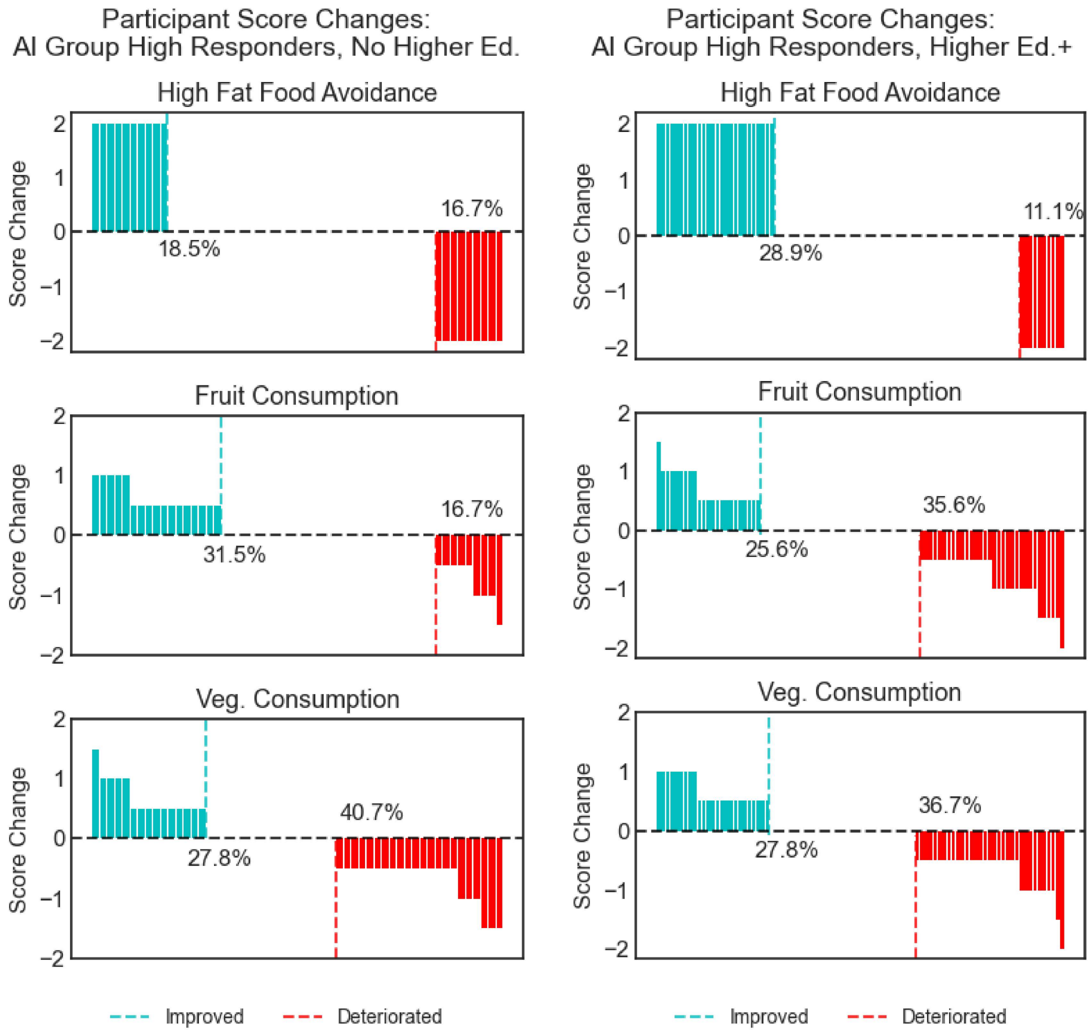

Figure 23: Nutrition score changes among AI intervention participants with and without some college education.

5.4.3 Age

Because type II diabetes is more likely to develop in midlife than in youth, the intervention was concerned with promoting behavior change among older participants [19]. Of similar importance was the development of sustainable healthy behaviors and greater health awareness among younger participants. Participants in the 18-35 age group were more likely to improve their cause and complication knowledge and high fat food avoidance, while participants in the 35+ group were more likely to improve their physical activity scores and vegetable consumption.

Cause and Complication Knowledge The 18-35 age group in the AI intervention exhibited greater rates of improvement and lower rates of deterioration in their type II diabetes cause and complication knowledge compared to the 35+ age group. While both groups benefit from being informed on type II diabetes, the significant improvements in knowledge scores among younger participants are meaningful, as these young adults now have a greater awareness and understanding of type II diabetes that they can draw on to lower their risk of type II diabetes over the course of their lifetimes.

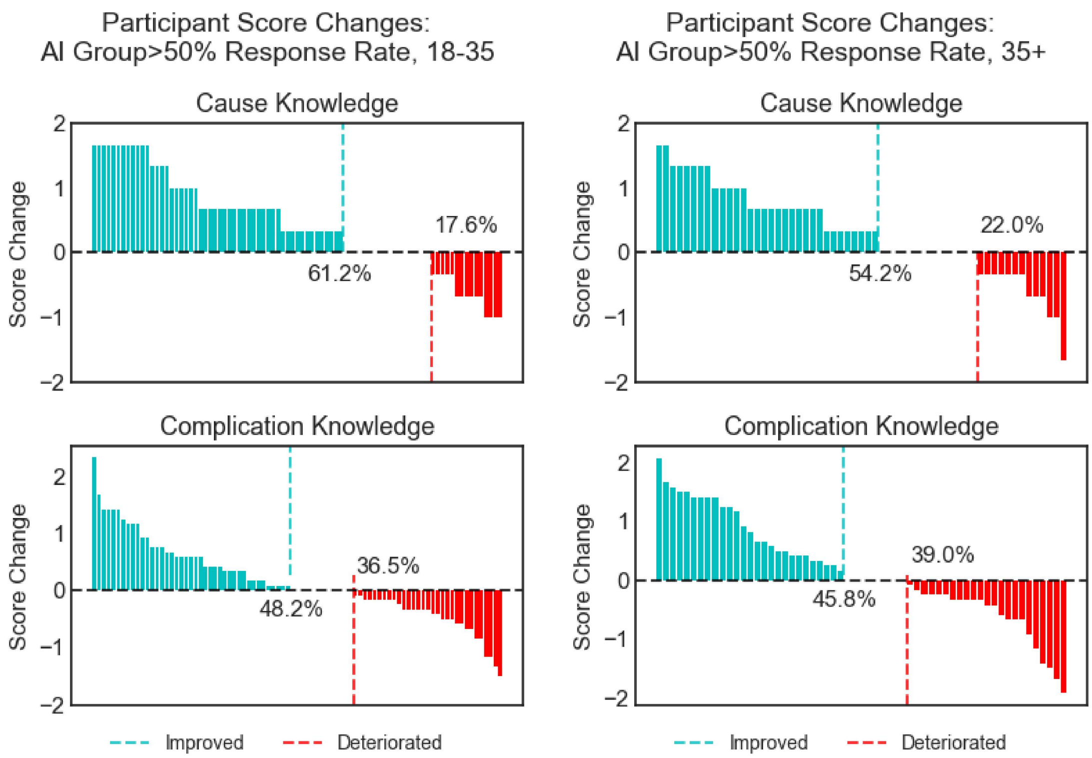

Figure 24: Knowledge score changes among older and younger AI intervention participants.

Physical Activity In all physical activity-related categories, the 35+ group was more likely to improve their scores. The older group was also slightly more likely to deteriorate in their daily average exercise times. Given the importance of physical activity in disease prevention and the fact that older adults have a higher risk of developing diabetes, the improvements in physical activity levels among older adults highlight the positive impact of the AI-based intervention.

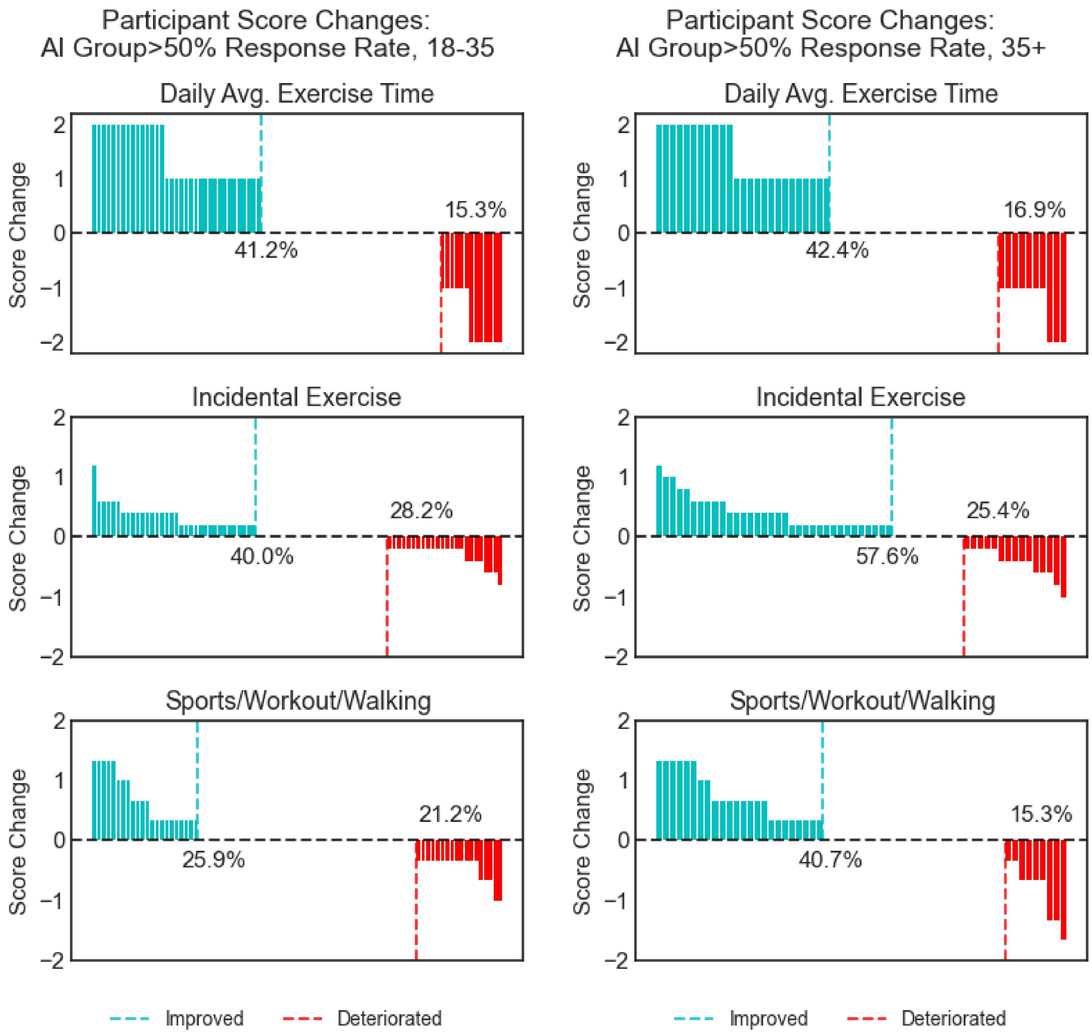

Figure 25: Physical activity score changes among older and younger AI intervention participants.

Nutrition The 18-35 group experienced higher rates of improvement and deterioration in their high fat food avoidance compared to the 35+ group, which saw more modest improvements alongside low rates of deterioration. The groups' changes in fruit consumption were similar, but the 35+ group demonstrated much higher rates of improvement in terms of vegetable consumption.

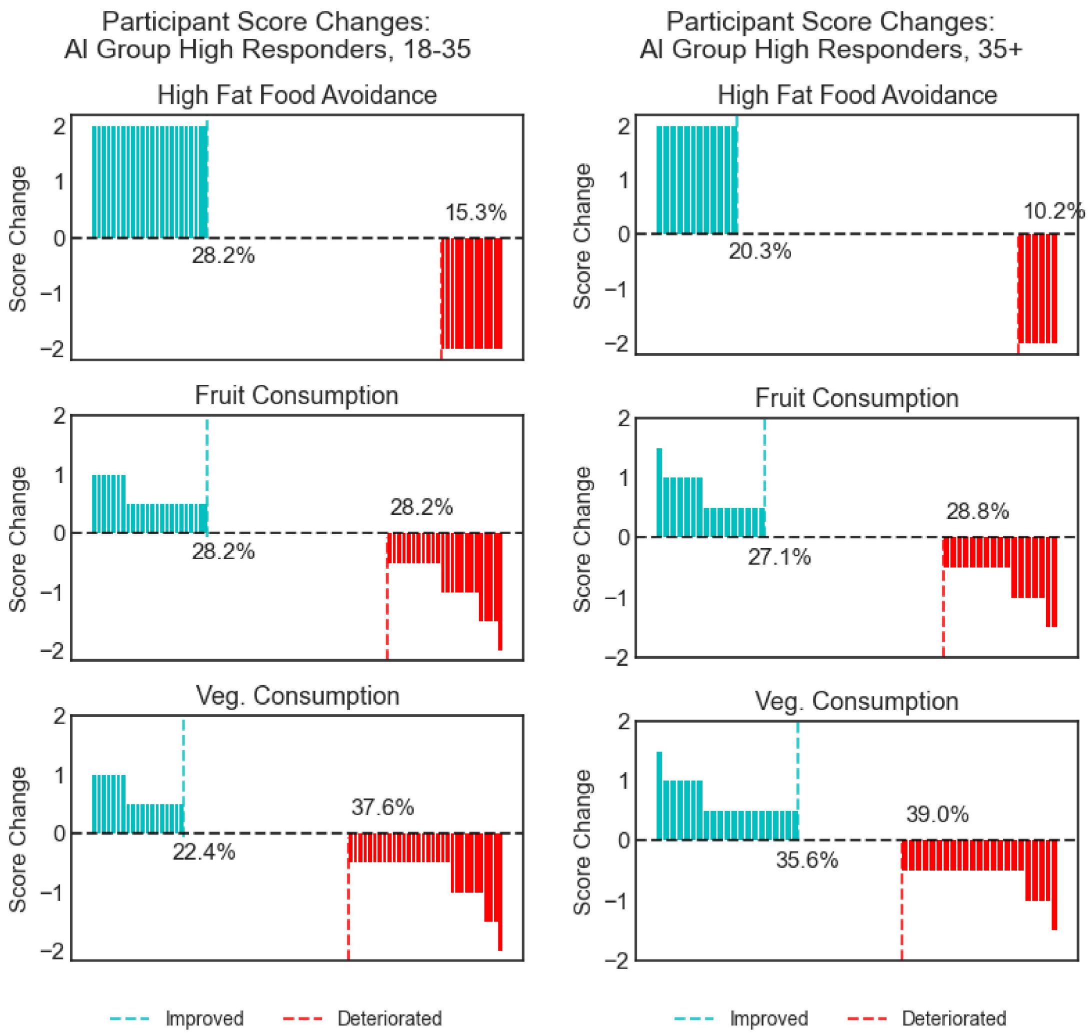

Figure 26: Nutrition score changes among older and younger AI intervention participants.

6 Discussion

The effectiveness of the AI-assisted intervention varied depending on participant engagement and demographic criteria. Still, when compared with the control group, highly engaged participants in the AI-based intervention exhibited notable behavior change in certain categories, specifically cause knowledge, daily average exercise time, incidental exercise, and high fat food avoidance. The success of the intervention in promoting physical activity is remarkable given that the original static intervention failed to provoke any significant level of improvement in physical activity.

As highlighted by the detailed results section, the amount of data produced by participants in the AI group provides the AI-based intervention with a major advantage over a traditional static intervention. Using the responses from the twice-weekly check-in, public health professionals can easily track longitudinal changes in participant behavior. Mid-intervention insights like these are more difficult to obtain when the intervention is not automated. Additional analysis, such as the demographic analyses in section five, are also easier to perform, enabling in-depth analysis of the impact of the intervention.

A key challenge for the AI-assisted intervention is maintaining participant engagement over the course of the study. As is typical for human behavioral experiments, participants dropped out of the AI-based intervention study, but many others simply stopped responding or rarely responded. Allowing participants to schedule the twice-weekly message and check-in at a convenient time may increase their participation in the intervention. While the AI-based intervention is effective in promoting behavior change in engaged participants, its utility is limited if participants are not motivated to stay involved with the intervention.

7 Appendix

7.1 Glossary

Machine Learning: A sub-field of artificial intelligence focused on teaching machines to uncover patterns and draw conclusions from data without the assistance of humans. Content algorithms are a widespread example of machine learning at work: Spotify, Netflix, YouTube, and other services all employ machine learning to examine user activity and recommend content that a user is likely to enjoy.

Machine Learning Algorithm: The method a machine uses to learn, or discover trends and relationships in input data. These methods are formulaic and often mathematically rigorous procedures that are trained to produce a machine learning model.

Training: The practice of teaching a machine learning algorithm to recognize patterns by running it on a practice dataset, known as training data. This process allows researchers to verify that the algorithm is making correct predictions on a dataset representative of the datasets the algorithm will process in the future.

Machine Learning Model: A trained algorithm capable of taking in an unfamiliar dataset and making predictions on the data based on the conclusions the algorithm drew from the training data.

Neural network: Machine learning model based on the human brain that is composed of layers of connected neurons and nonlinear functions that transform the original values to capture more information. At least one of the layers is hidden.

Hidden layer: Neural network layers that lay in between the input layer that received the original data and the output layer that reports predictions or results. Hidden layers apply nonlinear functions to the data to capture patterns in the data that are not linear.

Deep Neural Network: Neural network that contains more than one hidden layer. Deep neural networks (DNNs) may have dozens or hundreds of layers to improve performance. A drawback of models with more hidden layers is that they are less interpretable, meaning that we cannot trace exactly how the model arrived at its predictions. The DNN below has three hidden layers.

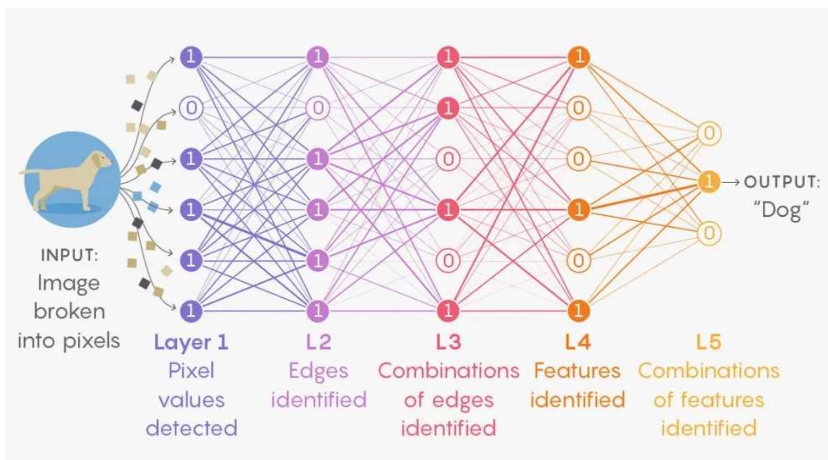

Figure 27: A deep neural network used for image recognition. [23]

Agent: An entity that learns how to make optimal decisions based on the outcomes of past decisions.

State: A collection of values that describe an agent's environment. For instance, if the agent is a rat learning to navigate a maze, the state is the position of the rat and the directions that the rat can move. An agent uses the state information to determine the optimal action in the current context.

Policy: The information that the agent uses to determine which action is best in the current environment state. The policy acts as a record of past actions and whether those actions produced positive, neutral, or negative outcomes. Continuing with the maze example, the policy is the rat's memory of its path through the maze; if the rat knows it has already taken a certain turn that led to a dead end and decides not to take that path again, it has used its policy to make an optimal decision.

Reinforcement Learning: A form of machine learning in which an agent must learning how to operate in an environment given little information. The agent takes actions and updates its policy depending on whether that action is rewarded. Through this process of trial and error, the agent learns which actions are optimal in different environment states.

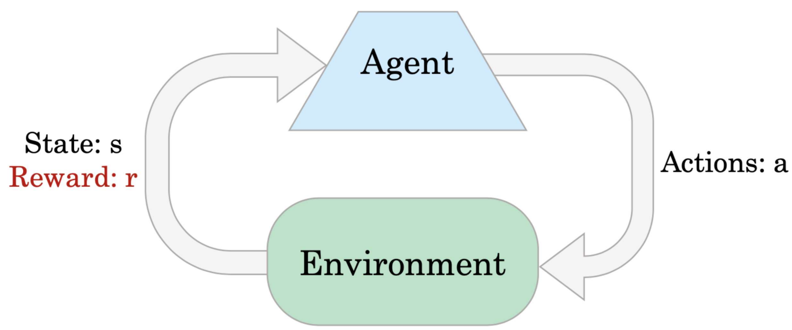

Figure 28: A conceptual representation of reinforcement learning.

Q-learning: A type of reinforcement learning in which the policy stores the overall reward of an action instead of the specific reward for that action, enabling the agent to make long-term predictions about optimal decisions. In each state, the agent predicts the maximum total reward from choosing a particular action and following that action to its conclusion. For example, if the rat knows that going right will take it closer to the end of the maze but will ultimately result in a dead end, that action will have a low overall reward. Under the Q-learning policy, the rat would choose to go left and take the long way to the exit since the overall reward is greater.

Deep Q-learning: A variant of Q-learning implemented with deep neural networks. Regular Q-learning uses a table that associates state/action pairs with the predicted reward value for choosing that particular action in that state. When there are many possible states and actions, storing data in a table is not possible, as it may grow impossibly large. In these cases, deep Q-learning can capture the state/action pairs and predicted values more efficiently than a table.

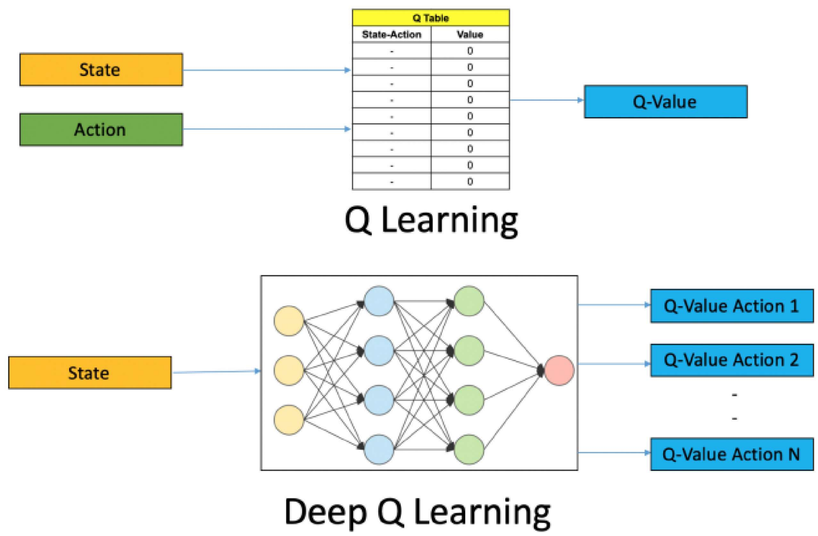

Figure 29: A comparison of a deep Q-learning model and a regular Q-learning model [1].

8 References

[1] "Deep Reinforcement Learning: Guide to Deep Q-Learning," Apr. 2021. [Online]. Available: https://www.mlq.ai/deep-reinforcement-learning-q-learning/

[2] "What are Recurrent Neural Networks?" https://www.ibm.com/cloud/learn/recurrent-neural-networks, Apr. 2021.

[3] W. Alpha, "Wolfram receiver operating characteristic curve demonstration," http://demonstrations.wolfram.com/HowReceiverOperatingCharacteristicCurvesWork/.

[4] A. Biswas, G. Aggarwal, P. Varakantham, and M. Tambe, "Learn to Intervene: An Adaptive Learning Policy for Restless Bandits in Application to Preventive Healthcare," *arXiv:2105.07965 [cs]*, July 2021, arXiv: 2105.07965. [Online]. Available: http://arxiv.org/abs/2105.07965

[5] A. Dautel, W. K. Härdle, S. Lessmann, and H.-V. Seow, "Forex exchange rate forecasting using deep recurrent neural networks," *Digital Finance*, vol. 2, 09 2020.

[6] G. Developers, "Machine learning glossary," https://developers.google.com/machine-learning/glossary.

[7] F. Gabbiani and S. J. Cox, *Mathematics for neuroscientists*. Academic Press, 2017.

[8] M. Hausknecht and P. Stone, "Deep Recurrent Q-Learning for Partially Observable MDPs," p. 9.

[9] K. J. Head, S. M. Noar, N. T. Iannarino, and N. Grant Harrington, "Efficacy of text messaging-based interventions for health promotion: A meta-analysis," *Social Science & Medicine*, vol. 97, pp. 41–48, Nov. 2013. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0277953613004474

[10] I. Hochberg, G. Feraru, M. Kozdoba, S. Mannor, M. Tennenholtz, and E. Yom-Tov, "A reinforcement learning system to encourage physical activity in diabetes patients," *ArXiv*, vol. abs/1605.04070, 2016.

[11] J. A. Killian, B. Wilder, A. Sharma, D. Shah, V. Choudhary, B. Dilkina, and M. Tambe, "Learning to Prescribe Interventions for Tuberculosis Patients Using Digital Adherence Data," *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 2430–2438, July 2019, arXiv: 1902.01506. [Online]. Available: http://arxiv.org/abs/1902.01506

[12] Lauffenburger, Julie C, Yom-Tov, Elad, Keller, Punam A, M. E. McDonnell, L. G. Bessette, C. P. Fontanet, E. S. Sears, E. Kim, K. Hanken, J. J. Buckley, R. A. Barlev, N. Haff, and N. K. Choudhry, "REinforcement learning to improve non-adherence for diabetes treatments by Optimising Response and Customising

Engagement (REINFORCE): study protocol of a pragmatic randomised trial," *BMJ Open*, vol. 11, no. 12, p. e052091, Dec. 2021. [Online]. Available: https://bmjopen.bmj.com/lookup/doi/10.1136/bmjopen-2021-052091

[13] Y. Li, J. Fan, Y. Wang, and K.-L. Tan, "Influence Maximization on Social Graphs: A Survey," *IEEE Transactions on Knowledge and Data Engineering*, vol. 30, no. 10, pp. 1852–1872, Oct. 2018, conference Name: IEEE Transactions on Knowledge and Data Engineering.

[14] Y. Lu, D. Wang, Q. Meng, and P. Chen, "Towards Interpretable Deep Learning Models for Knowledge Tracing," *arXiv:2005.06139 [cs]*, May 2020, arXiv: 2005.06139. [Online]. Available: http://arxiv.org/abs/2005.06139

[15] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, Feb. 2015. [Online]. Available: http://www.nature.com/articles/nature14236

[16] D. Mohan, J. J. H. Bashingwa, N. Tiffin, D. Dhar, N. Mulder, A. George, and A. E. LeFevre, "Does having a mobile phone matter? Linking phone access among women to health in India: An exploratory analysis of the National Family Health Survey," *PLOS ONE*, vol. 15, no. 7, pp. 1–16, July 2020, publisher: Public Library of Science. [Online]. Available: https://doi.org/10.1371/journal.pone.0236078

[17] R. Nyrup, "From General Principles to Procedural Values: Responsible Digital Health Meets Public Health Ethics," *Frontiers in Digital Health*, vol. 3, p. 690417, July 2021. [Online]. Available: https://www.frontiersin.org/articles/10.3389/fdgth.2021.690417/full

[18] A. Pfammatter, B. Spring, N. Saligram, R. Davé, A. Gowda, L. Blais, M. Arora, H. Ranjani, O. Ganda, D. Hedeker, S. Reddy, and S. Ramalingam, "mHealth Intervention to Improve Diabetes Risk Behaviors in India: A Prospective, Parallel Group Cohort Study," *Journal of Medical Internet Research*, vol. 18, no. 8, p. e207, Aug. 2016. [Online]. Available: http://www.jmir.org/2016/8/e207/

[19] R. Pradeepa and V. Mohan, "Epidemiology of type 2 diabetes in india," *Indian Journal of Ophthalmology*, vol. 69, no. 11, p. 2932, 2021.

[20] S. Ranganathan, K. Nakai, and C. Schonbach, *Encyclopedia of Bioinformatics and Computational Biology: ABC of Bioinformatics*. Elsevier, 2018.

[21] S. Umbrello and I. van de Poel, "Mapping value sensitive design onto AI for social good principles," *AI and Ethics*, vol. 1, no. 3, pp. 283–296, Aug. 2021. [Online]. Available: https://link.springer.com/10.1007/s43681-021-00038-3

[22] J. Weltz, A. Volfovsky, and E. B. Laber, "Reinforcement Learning Methods in Public Health," *Clinical Therapeutics*, vol. 44, no. 1, pp. 139–154, Jan. 2022. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0149291821004550

[23] N. Wolchover, "New Theory Cracks Open the Black Box of Deep Neural Networks," *Wired*, section: tags. [Online]. Available: https://www.quantamagazine.org/new-theory-cracks-open-the-black-box-of-deep-learning-20170921/

[24] E. Yom-Tov, G. Feraru, M. Kozdoba, S. Mannor, M. Tennenholtz, and I. Hochberg, "Encouraging physical activity in patients with diabetes: Intervention using a reinforcement learning system," *Journal of Medical Internet Research*, vol. 19, no. 10, p. e338, oct 2017. [Online]. Available: https://doi.org/10.2196%2Fjmir.7994

[25] C. Yu, J. Liu, and S. Nemati, "Reinforcement learning in healthcare: A survey," *ACM Computing Surveys (CSUR)*, vol. 55, pp. 1 – 36, 2019.