Exploring the Influence of Multiple Virtual Agents for Addressing Barriers to Cancer Clinical Trial Participation

Rashi Ghosh
rashighosh@ufl.edu
University of Florida
Gainesville, Florida, USA

Christopher You
christopheryou@ufl.edu
University of Florida
Gainesville, Florida, USA

Roshan Venkatakrishnan
rvenkatakrishnan@ufl.edu
University of Florida
Gainesville, Florida, USA

Benjamin Lok
lok@ufl.edu
University of Florida
Gainesville, Florida, USA

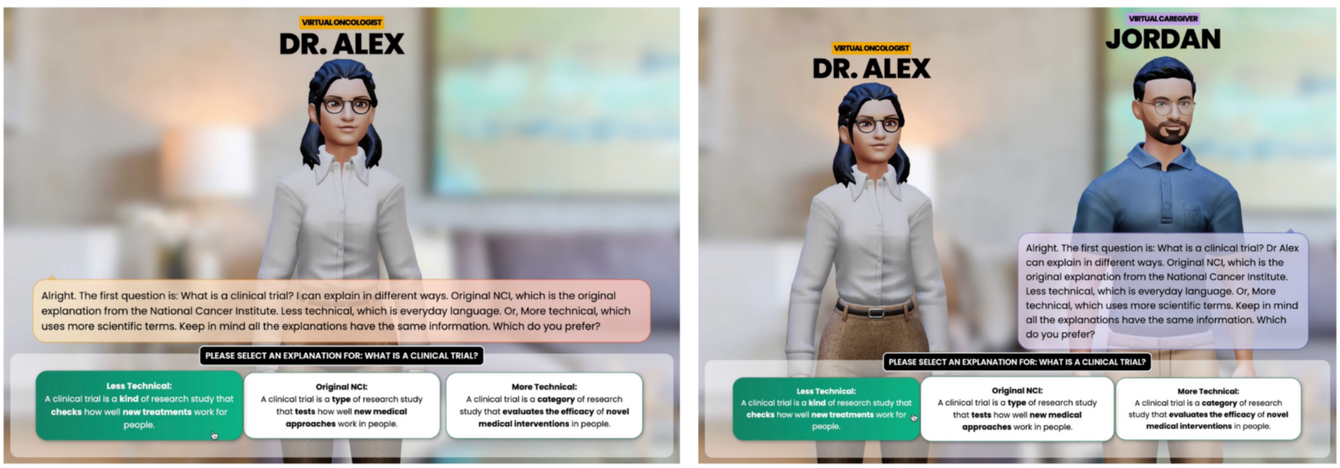

Figure 1: The single virtual agent (left) and multiple virtual agent (right) interfaces

Abstract

Clinical trials are essential for advancing cancer treatment, yet participation remains low due to barriers such as limited knowledge and negative attitudes, largely due to poor communication. Virtual agents have shown promise in delivering health information, and multi-agent systems may better support complex, human-like interactions. However, their potential for addressing clinical trial barriers remains underexplored. In this work, we conducted a between-participants study with 61 users comparing a single-agent intervention with a multi-agent intervention. Our findings show that multiple agents can improve knowledge outcomes. However, their impact on behavioral intention varies based on users' health literacy. This work enhances our understanding of leveraging multi-agent interventions to support effective health communication and highlights the importance of user health literacy when designing virtual agent interventions. Our findings offer early insights into the potential of using multiple virtual agents to address real-world health challenges, particularly barriers to cancer clinical trial participation.

Conference acronym 'XX, Woodstock, NY

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

CCS Concepts

• **Human-centered computing → User studies**.

Keywords

virtual agents, multi-agent systems, healthcare interventions, clinical trials, health literacy

ACM Reference Format:
Rashi Ghosh, Christopher You, Roshan Venkatakrishnan, and Benjamin Lok. 2018. Exploring the Influence of Multiple Virtual Agents for Addressing Barriers to Cancer Clinical Trial Participation. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX)*. ACM, New York, NY, USA, 10 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

Clinical trials are critical to the development of new cancer treatments. Yet, adult participation in cancer clinical trials remains low – typically between 2% and 8% [57, 58]. A range of barriers contribute to this under-enrollment, including limited awareness and poor communication about clinical trials [30, 33, 56]. Poor communication often results from failing to address patients' information needs using appropriate language and health literacy levels [22]. This is particularly critical for individuals with low health literacy, who may struggle to interpret complex health information and are therefore less likely to make informed decisions about clinical trial participation [28]. In such cases, having another person involved in the conversation in addition to the doctor can help clarify information and support understanding [31]. In cancer care literature, patients are often accompanied by a third person during their doctor appointments [27], which has been associated with improved patient engagement and understanding [54]. For example, patients may feel uncomfortable asking for clarification when receiving information from a doctor, and companions can play a key role in ensuring the information is delivered in a way that makes sense to the patient [11, 27, 45]. Hence, this third person can offer informational support and help facilitate communication [27], leading to richer information exchange and better comprehension [25].

This real-world context provides an opportunity to explore a novel approach to expand accessibility in digital healthcare interventions using multiple virtual agents. Virtual agents are often implemented as computer-generated characters that simulate human-like behavior [59], enabling people to ascribe real-world social roles such as doctors and health coaches to virtual agents [40, 50] to deliver health information. However, existing literature has largely explored interventions with one virtual agent (single-agent interventions) and demonstrates mixed outcomes. Single-agent interventions have been shown to effectively communicate complex health information for people of varying information needs [6, 64], including people with low health literacy [7]. At the same time, interventions using a single agent have been found to be difficult to interact with for lower health literacy users [14] while providing better comprehension-related benefits for higher health literacy users [8].

Therefore, while virtual agents are a promising avenue for communicating health information, single-agent interventions may not sufficiently support users of varying health literacy in contexts such as addressing clinical trial barriers. Drawing from real-world cancer conversations where a third-person companion aids with information comprehension [27, 29], a virtual agent intervention may provide better outcomes when using more than one virtual agent (multi-agent intervention). Indeed, prior work highlights multiple agents' potential in healthcare contexts to provide a richer and more realistic communication experience, such as delivering tailored health coaching styles [5] and presenting complementary health information without contradiction [24] Hence, a multi-agent intervention may offer a richer platform more closely resembling information delivery during real-world cancer consultations by simulating a three-way conversation between a user and two virtual agents (doctor and companion).

To explore the effectiveness of a multi-agent intervention for addressing barriers to clinical trial participation, we conducted a between-participants study comparing a single-agent intervention (doctor) to a multi-agent intervention (doctor and companion). We explore the following research question: ***How does a multi-agent intervention impact knowledge, attitudes, and intentions regarding cancer clinical trial participation?*** Our hypotheses are as follows:

- ***H1:*** *A multi-agent intervention can improve individuals' knowledge about clinical trials.*
- ***H2:*** *A multi-agent intervention can improve individuals' attitudes towards clinical trials.*
- ***H3:*** *A multi-agent intervention can improve individuals' intention to participate in clinical trials.*

In light of evidence that health literacy can influence how individuals process health information, we tested our hypotheses exploring the role of health literacy in this context. While health literacy is not included in our primary hypotheses, we sought to investigate its potential influence given its theorized relevance to clinical trial communication. [22, 28]

2 Related Works

In this section, we provide an overview of how virtual agents have been used to communicate health information and address clinical trial barriers. We also discuss prior work with multi-party interventions.

2.1 Virtual Agents to Address Information Comprehension and Health Literacy

Prior work shows that interventions using a single virtual agent are an acceptable and useful medium for health interventions across all levels of health literacy [6]. Virtual agents can deliver information based on user traits in ways that meet the user's information needs, leading to improvement in cancer-related knowledge topics [64]. Low health literacy users have even reported high levels of satisfaction and preferences for information delivered by virtual agents over humans [7, 8]. However, prior work has also found contrasting effects. Users with lower health literacy have been shown to have more difficulty interacting with adaptive virtual agents [14], while users with higher health literacy benefited more in terms of comprehension with virtual agents [8].

Therefore, it is unclear how to design virtual agent interventions to be accessible for users with different levels of health literacy. This is highlighted by conflicting results regarding how single-agent interventions can deliver information to support people of different health literacies. Therefore, this work contributes to virtual agent design by exploring whether the addition of a second virtual agent can improve user outcomes regarding barriers to clinical trial participation, for all levels of user health literacy.

2.2 Virtual Agents to Address Clinical Trial Barriers

Previous research has demonstrated that virtual agent interventions are a promising strategy to address clinical trial barriers [9, 10, 38]. Mozgai et al's initial feasibility study suggests that virtual agents can be used to address barriers to clinical trials and even promote intention to enroll [38]. Additionally, literature shows positive attitudes towards virtual agents addressing clinical trial barriers if their appearance and communication strategy are appropriately tailored [10]. Finally, virtual agents can provide a more comprehensible way to access clinical trial information [9], particularly for users with lower health literacy.

However, previous research has mainly explored attitudes and perceptions of virtual agents in this space, and it is still unclear if virtual agents can actually improve patient outcomes after addressing clinical trial barriers. Therefore, our work contributes to our understanding of virtual agent interventions by evaluating their impact on users' knowledge, attitudes, and intentions toward clinical trials.

2.3 Multiple Virtual Agents

Although most prior work in healthcare virtual agent interventions has focused primarily on dyadic interactions between a single virtual agent and a user [53, 62, 63], there has been some work exploring the use of multiple virtual agents. Kantharaju et al. found that using two virtual agents instead of one was more successful at persuading users of movie ratings [23]. This was true not only when both virtual agents were directly interacting with the user but also when the two virtual agents were interacting with each other, with the user simply observing. Multiple agents have also been explored in health-related contexts to deliver health information from a variety of viewpoints. Using multiple agents could allow each virtual agent to have their own expertise, providing an opportunity to present multiple perspectives without the agents contradicting themselves [24]. For example, Beinema et al. used multiple agents with different coaching styles tailored to a user's motivations to promote healthy living [5]. However, prior work has also demonstrated that a single agent may be better suited for health behavior change than multiple agents, although this work used physical robots instead of virtual agents [1].

Given the limited prior work on the effectiveness of multiple agents in health-related interventions, our work sought to expand our understanding of the use of multiple agents. Our work contributes to virtual agent design by evaluating how multiple virtual agents can influence user outcomes regarding clinical trial barriers.

3 Virtual Agent Intervention Design

This section describes the design of the virtual agent intervention in order to explore our research question: ***How does a multi-agent intervention impact knowledge, attitudes, and intentions regarding cancer clinical trial participation?*** Since our study aims to address breast cancer patients' understanding of clinical trials, we use the term *intervention* in line with healthcare literature, where intervention may refer to educational efforts like providing information to enhance understanding [51]. First, we describe how the virtual agents and web-based platform were created, followed by how the dialogue and conversation were implemented. Then, we provide a detailed breakdown of each part of the virtual agent intervention.

3.1 Agent and System Design

Virtual Agent Roles. Because this study's motivation draws from real-world companions facilitating information delivery between a doctor and a patient [27], we developed two systems: (1) a single-agent system featuring a virtual doctor, Dr. Alex, and (2) a multi-agent system featuring Dr. Alex *and* a virtual companion, Jordan. The same virtual doctor was used in both the single- and multi-agent conditions. The doctor role was presented as an oncologist because they are often the ones to bring up clinical trials to patients [3, 12, 15]. The companion role was presented as a caregiver because they often accompany patients to appointments to provide support and help process information [4, 13].

The function of the companion virtual agent in the multi-agent condition was to facilitate communication between the virtual doctor and the user. Patients often feel hesitant to seek clarification from doctors, and companions can help ensure information is delivered in a patient-centered way [27]. Similarly, the virtual companion prompted the virtual doctor to explain information in ways aligned with users' preferences [11, 45] (see Section 3.2 for details).

3D Virtual Agent System. The virtual agents were developed using Ready Player Me [49], a web-based avatar creation platform for building and customizing 3D avatars, and displayed waist up during the intervention. The virtual agents were integrated into a NodeJS application and deployed to the web via Amazon Web Services. The virtual agents' voice and nonverbal behaviors were implemented using *TalkingHead* [36] , an open-source repository that lip-syncs and animates virtual characters using ThreeJS. Google's text-to-speech engine was used to generate speech audio [19]: en-US-Neural2-C was used for the virtual doctor, and en-US-Neural2-J was used for the virtual caregiver. Google text-to-speech provides audio and individual word time stamps, which *TalkingHead* uses to generate new visemes, timestamps, and word durations for real-time lip-sync animations. Additionally, the *TalkingHead* repository provides templates for nonverbal behaviors such as poses, gestures, and animations. For this study, we used default settings for nonverbal behavior, which automatically applies behaviors such as *waving* when the characters first introduce themselves, *idle* animations, and minor head nods and hand movements when *speaking*.

Dialogue and Response Generation. The virtual agents' dialogue were pre-scripted and adapted from an existing web-based intervention for addressing barriers to clinical trial participation [35]. The intervention was centered around answering pre-defined questions from a subset of 27 questions regarding clinical trials from Meropol et al [35]. Although the conversation was pre-scripted in order to ensure greater experimental control across single- and multi-agent conditions, we used AI to pre-generate all responses. This was done to align our methodology with emerging practices in developing virtual agent interventions, which are increasingly using large language models to generate dialogue for more real-time, adaptive conversations [52].

Because prior work suggests people differ in how much complexity they prefer when receiving information [28], we pre-generated three responses of varying language complexity for each question for users to choose from. In the multi-agent condition, the companion virtual agent presented these options, reflecting its role in

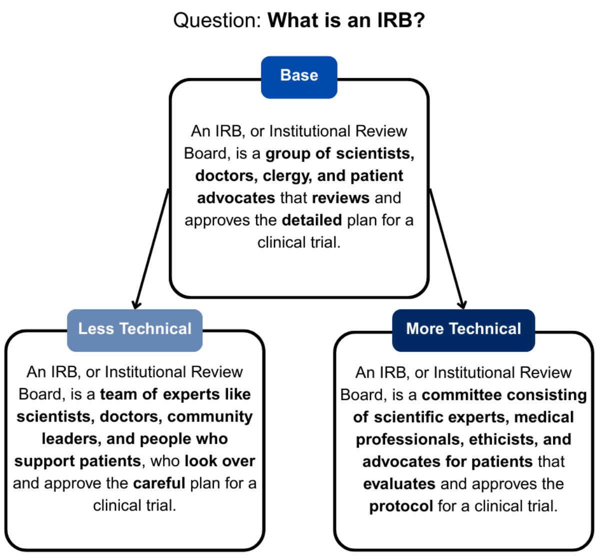

Figure 2: Examples of explanation preferences for the question *"What is an IRB?"* Specific words were bolded to emphasize the language literacy differences.

helping facilitate information delivery (see Section 3.2 for more details). The three response types were as follows:

- *Base* - Information directly from the National Cancer Institute (NCI) [41] with no modification
- *Less Technical* - Modifications to the *base* response using simpler language and less medical jargon based on plain language guidelines [46], which is the recommended practice for communicating lower health literacy [28, 44]
- *More Technical* - Modifications to the *base* response using more complex language, medical jargon, and technical terms

All responses were generated using OpenAI's Assistants API [43], starting with the *base* response. For each question, the *base* response was generated using a vector store, which served as a knowledge base with information from the National Cancer Institute (NCI) [41]. This source is reputable and commonly used in existing cancer clinical trial-related interventions [35]. The *less technical* and *more technical* responses were constructed from the *base* response. The content of the three response types was the same – only the language differed (Figure 2). The *base, less technical*, and *more technical* responses were validated by health communication experts. A total of 81 responses were generated: three versions of varying language complexity for each of the 27 questions. All 81 responses were pre-generated to ensure consistency. Each response was around 75 words in length.

3.2 Virtual Agent Intervention

The virtual agent intervention was fully web-based. In the single-agent condition, participants only saw Dr. Alex. In the multi-agent condition, participants saw both Dr. Alex and Jordan at the same time. Participants interacted with the virtual agents using multiple-choice button options (Figure 3), similar to other virtual agent studies for health-related interventions [34] [61]. The intervention structure consisted of the following steps:

- *Introduction.* The virtual agents introduced themselves and explained their role in the intervention. In the single-agent condition, only Dr. Alex introduced themself. In the multi-agent condition, Dr. Alex introduced themself first, followed by Jordan. Dr. Alex's introduction was kept consistent between both conditions.
- *Explanation Preference Phase.* This phase was designed for participants to explore the three explanation types (see Section 3.1 for more details) that the virtual agent could use to explain clinical trial concepts. The virtual agent(s) guided the participants through three questions (*What is a clinical trial*, *What is an IRB*, and *What is informed consent*? [35]). For each question, participants were shown three explanation options with key language differences bolded and were asked to select which option they preferred: original NCI (base), less technical, or more technical (Figure 1). In the single-agent condition, Dr. Alex presented the explanation options and responded based on the user's selection. In the multi-agent response, Jordan presented the explanation options, and Dr. Alex responded based on the user's selection. After presenting the three introductory questions and exploring the explanation options, a preferred explanation type was determined. This was used consistently in the subsequent intervention phase. Participants had the opportunity to change their explanation preference at any point during the remainder of the intervention, however, none of the participants chose to do so.
- *Question-Answering Phase.* The purpose of this phase was to answer relevant questions about clinical trials using the participant's explanation preference. Participants were required to ask a total of seven questions using button input. These questions were pre-selected from a list of 27 pre-defined questions based on the user's responses to the clinical trials knowledge and attitudes questionnaires from the pre-survey [35]. For each question, Dr. Alex responded using the explanation preference the user previously selected (for both conditions). At any point, the user could click on the virtual agent (in the single-agent condition, Dr. Alex; in the multi-agent condition, Jordan) to change their explanation preference.
- *Conclusion.* After the user asked all seven questions, the virtual agent concluded the intervention by instructing the participant to click a button to complete the post-survey (in the single-agent condition, Dr. Alex; in the multi-agent condition, Jordan).

4 Methodology

To investigate the impact of multiple virtual agents in an intervention designed to address barriers to cancer clinical trial participation, we conducted an online, asynchronous between-participants study. We manipulated the number of virtual agents across two experimental conditions: **single-agent** and **multi-agent**. In the single-agent condition, participants interacted with only one virtual

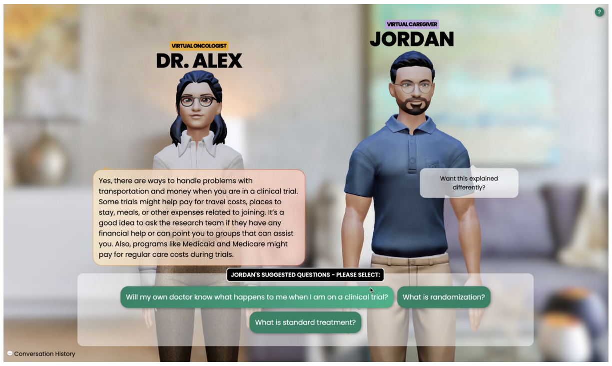

Figure 3: Multiple-choice button input interface during the question-asking phase for the multi-agent intervention

agent throughout the entire intervention. In the multi-agent condition, participants interacted with two virtual agents. The content and dialogue were kept consistent between the two conditions.

4.1 Participants

Participants were recruited from Prolific, an online research recruitment platform [48]. The study was approved by the university's IRB prior to data collection (IRB#ET00045896). As part of the study's inclusion criteria, participants had to be 18 years or older, reside in the United States, and have a current or past cancer diagnosis. Monetary compensation of 5.00 USD was provided for participation. A total of 61 participants were recruited, all of whom completed the entire study.

4.2 Measures

Health Literacy. Participants' self-reported health literacy was measured in the pre-survey using the BRIEF Health Literacy Screening Tool [21]. All four items (see Table 1) were summed for a total score. Participants were grouped as either *inadequate* (summed score is 4-16) or *adequate* (summed score is 17-20) health literacy based on their summed scores, which is common practice in health communication literature [20] [8].

Health Literacy Item	Anchors
How often do you have someone help you read hospital materials?	(1 = Always, 3 = Sometimes, 5 = Never)
How confident are you filling out medical forms by yourself?	(1 = Not at all, 3 = Somewhat, 5 = Extremely)
How often do you have problems learning about your medical condition because of difficulty understanding written information?	(1 = Always, 3 = Sometimes, 5 = Never)
How often do you have a problem understanding what is told to you about your medical condition?	(1 = Always, 3 = Sometimes, 5 = Never)

Table 1: Self-Reported Health Literacy items from the BRIEF Health Literacy Screening Tool [21]

Knowledge About Clinical Trials. Participants' knowledge about clinical trials was assessed in the pre- and post-survey. The knowledge questionnaire included 19 true-false items adopted from [35] and was scored as a sum of total correct.

Attitudes Toward Clinical Trials. Participants' attitudes towards clinical trials was measured in the pre- and post-survey with a 15-item questionnaire adopted from [35]. Participants used a 0-100 slider scale to indicate their level of agreement with each item, where 0 = "Strongly disagree," 50 = "Neither agree nor disagree," and 100 = "Strongly agree." Overall attitudes towards clinical trials were scored as an average of all 15 items.

Intention to Participate in Clinical Trials. Participants' intention towards participating in clinical trials was measured in the pre- and post-survey with a single item, as seen in other research on clinical trial barriers [17, 37, 47]. Participants used a 0-100 slider scale to indicate their level of agreement with the following item: *"I intend to participate in a clinical trial in the future,"* where 0 = "Strongly disagree," 50 = "Neither agree nor disagree," and 100 = "Strongly agree."

Perceptions of the Virtual Agent. Participants' perceptions of the virtual agent were measured in the post-survey using two single-item questions adopted from [63] (see Table 2). In the multi-agent condition, participants answered each question for each virtual agent.

Perceptions of the Virtual Agent	Anchors
How much did you trust the virtual character?	(0 = Not at all, 100 = Very much)
How knowledgeable was the virtual character?	(1 = Not at all, 100 = Very knowledgable)

Table 2: Single-item measures assessing perceptions of the virtual agents [63]

4.3 Procedure

First, participants completed the pre-survey on Qualtrics, beginning with informed consent. Then, participants completed questionnaire items on knowledge about clinical trials, attitudes toward clinical trials, intent to participate in a clinical trial, and self-reported health literacy. After completing the pre-survey, participants were randomized to either the *single-agent* or *multi-agent* condition. They were redirected to the intervention website and completed the process described in Section 3.2. After completing the entirety of the virtual agent intervention, participants were redirected to Qualtrics where they proceeded to complete the post-survey. Participants completed the same questionnaire items from the pre-survey on knowledge, attitudes, and intentions, as well as demographics. An overview of the entire procedure is outlined in Figure 4. Upon completion, participants were thanked and provided monetary compensation for their time. On average, it took about 25 minutes to complete the entire procedure (pre-survey, intervention, and post-survey).

5 Results

Data analysis was conducted using R. Due to violation of normality assumptions assessed by Shaprio-Wilk tests, we used the Align Rank

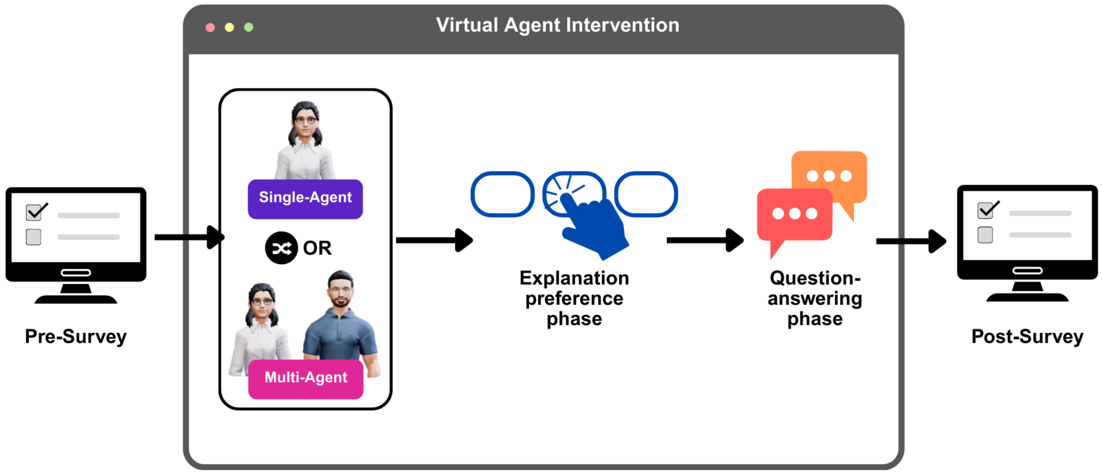

Figure 4: Overview of the study procedure.

Transform (ART) ANOVA for nonparametric factorial analysis [60], similar to other work [16, 34]. Our primary analyses focused on assessing the main effect of intervention condition (single-agent versus multi-agent) on outcome variables. In addition, we included analyses to explore whether participant health literacy (adequate versus inadequate) moderated the effect of intervention condition on these outcomes. To do so, health literacy and its interaction with condition were included as predictors in the ART ANOVA models.

To assess intervention outcomes, we used difference scores calculated by subtracting pre-intervention values from post-intervention values, as seen in similar work evaluating pre-post scores [64]. Significant effects are reported at the conventional threshold of $p < 0.05$, with Tukey HSD adjustments applied for post hoc comparisons to control for family-wise error. Following prior HCI work, we also report trending effects for results with $0.05 \leq p < 0.10$ [18, 42, 53, 55] to indicate patterns that may warrant further investigation, while interpreting them with caution. Finally, we report effect sizes as follows: small effect size $0.01 \leq \eta p^2 < 0.06$, medium effect size $0.06 \leq \eta p^2 < 0.14$, and high effect size $\eta p^2 \geq 0.14$ [32].

To analyze categorical explanation preference choices, we used chi-square goodness-of-fit tests to assess whether selections within subgroups deviated from uniform distribution. Where expected cell counts were below threshold assumptions (e.g., observed counts of zero), we used multinomial exact tests to ensure statistical validity.

5.1 Participants

A total of 61 participants were recruited, all of whom completed the entire study. Fifty-seven (n=57) participants reported their age: 20-73 years old (median=49). Fifty-eight (n=58) participants reported their gender with 68.9% female (n=42) and 26.2% (n=16) male. Fifty-five (n=55) reported their race/ethnicity as follows: 1.6% Asian (n=1), 9.8% Black/African American (n=6), 3.3% Hispanic/Latin American (n=2), 73.8% White - Non Hispanic/Non Latin American (n=45), and 1.6% Not Listed (n=1). Of the 61 total participants, 52.5% (n=32) completed the *single-agent* condition and 47.5% (n=29) completed the *multi-agent* condition. Regarding health literacy, 63.9% (n=39) participants had *adequate* health literacy and 36.1% (n=22) had *inadequate* health literacy.

No significant difference was observed for participants between the two conditions regarding gender and mean age on knowledge, attitude, or intention outcomes. Additionally, baseline analyses indicated no significant imbalances for pre-scores for knowledge, attitudes, and intention toward clinical trials between the two conditions.

5.2 Impact on Intervention Outcomes

Knowledge. The ART ANOVA revealed a significant main effect of condition ($F(1, 57) = 5.27, p = .025, \eta p^2 = 0.08$), indicating a medium effect size (Figure 5). Participants in the multi-agent condition had significantly higher knowledge increase compared to participants in the single-agent condition. There were no main effects of health literacy or interaction effects between condition and health literacy.

Attitudes. The ART ANOVA revealed no main effects of condition or health literacy on pre-post attitude change. There were also no interaction effects between condition and health literacy.

Intention. The ART ANOVA revealed no main effects of condition or health literacy on pre-post intention change. However, there was a significant interaction effect between condition and health literacy ($F(1, 57) = 6.03, p = .017, \eta p^2 = 0.10$), indicating a medium effect size. Post hoc comparisons indicated a trend suggesting that participants with adequate health literacy experienced greater improvements in intention than those with inadequate health literacy within the multi-agent condition.

5.3 Perceptions of Virtual Agents

Trust. The ART ANOVA revealed no main effects of condition. However, there was a significant main effect of health literacy ($F(1, 57) = 6.70, p = .012, \eta p^2 = 0.11$), indicating a medium effect size. Additionally, there was a significant interaction effect between condition and health literacy ($F(1, 57) = 5.53, p = .022, \eta p^2 =$

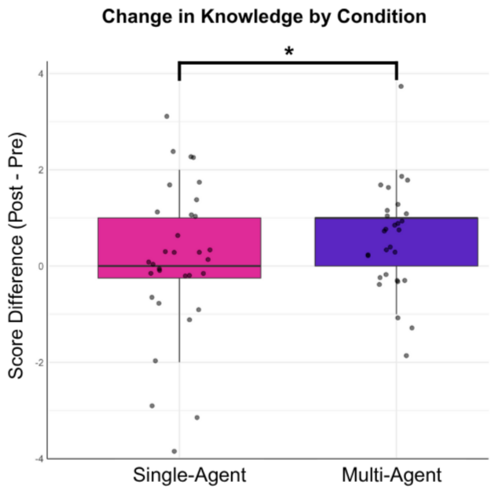

Figure 5: Participants' knowledge gain by condition. Significance codes: (*) denotes $p < 0.05$.

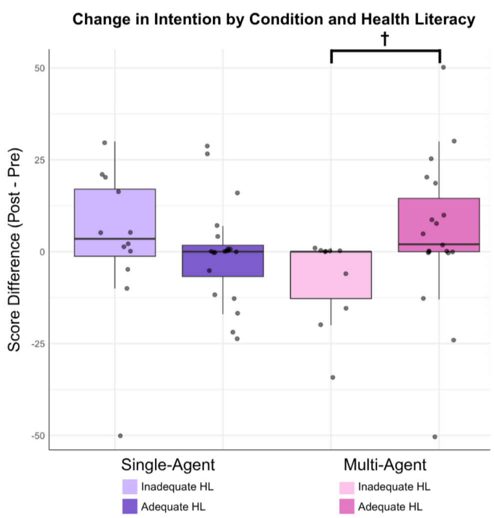

Figure 6: Participants' intention improvement by condition and health literacy. Significance codes: (†) denotes $p < 0.1$.

0.09), indicating a medium effect size. Specifically, in the multi-agent condition, participants with adequate health literacy had significantly higher trust compared to participants with inadequate health literacy ($t(57) = 3.53, p = .005$) ((Figure 7). Trends suggested that participants in the multi-agent condition with adequate health literacy also had higher trust than participants in the single-agent condition with both adequate ($t(57) = 2.44, p = .081$) and inadequate health literacy ($t(57) = 2.39, p = .090$).

Knowledgable. The ART ANOVA revealed no main effects of condition. However, there was a significant main effect of health literacy ($F(1, 57) = 5.02, p = .029, \eta p^2 = 0.08$), indicating a medium effect size. Specifically, participants with adequate health literacy perceived the virtual agents as significantly more knowledgeable

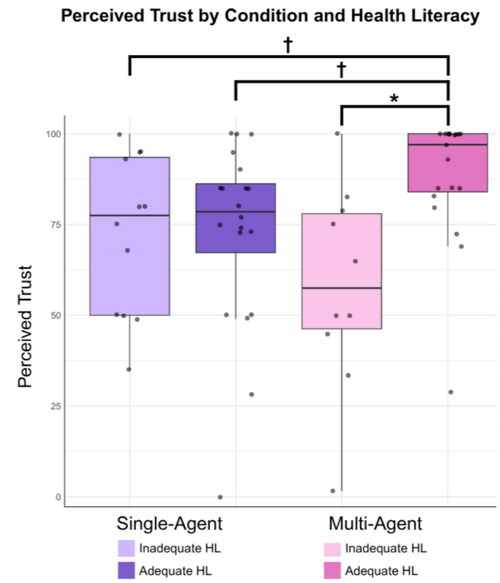

Figure 7: Participants' perceived trust of the virtual gents by condition and health literacy. Significance codes: (*) denotes $p < 0.05$, (†) denotes $p < 0.1$..

compared to participants with inadequate health literacy ($t(57) = 2.24, p = .029$). There were no interaction effects between condition and health literacy.

5.4 Selection of Explanation Preference

To examine whether explanation preferences (less technical, base, more technical) were distributed differently across health literacy and condition, a series of chi-square goodness-of-fit tests were conducted. No significant deviations from equal preference distribution were observed among participants with adequate health literacy in either the single-agent condition or the multi-agent condition. Similarly, inadequate health literacy participants in the single-agent condition showed no significant preference pattern. However, among participants with inadequate health literacy, the overall distribution of explanation preferences significantly deviated from uniform expectations ($\chi^2(2) = 10.18, p = 0.006$), revealing that less technical explanations were selected most frequently (Figure 8). This effect was especially pronounced for this group in the multi-agent condition ($\chi^2(2) = 7.40, p = .025$). A multinomial exact test confirmed this result ($p = 0.022$), and post hoc standardized residuals suggested that the less technical was selected more than expected ($z = 2.01$), while more technical was never selected ($z = -1.83$). A follow-up Fisher's Exact Test comparing the less technical explanation preference between conditions for inadequate health literacy participants revealed no significant difference.

6 Discussion

Our work explored how a multi-agent intervention addressing clinical trial barriers could impact user outcomes. In this section, we

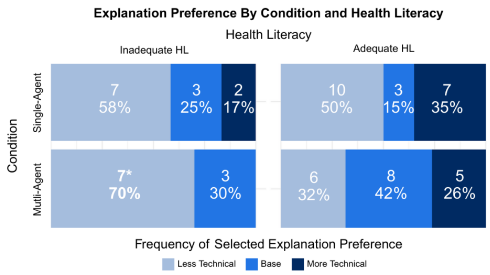

Figure 8: Frequency of selected explanation preference by condition and health literacy. Each bar segment shows the number and percentage of participants within each health literacy group who selected each explanation type in each condition. Asterisks (*) indicate statistically significant deviations from uniform choice based on chi-square tests.

discuss our results for our hypotheses: ***(H1)*** *A multi-agent intervention can improve individuals' knowledge about clinical trials*, ***(H2)*** *A multi-agent intervention can improve individuals' attitudes towards clinical trials*, and ***(H3)*** *A multi-agent intervention can improve individuals' intention to participate in clinical trials.*

Our findings support *H1*, demonstrating that participants in the multi-agent condition experienced significantly greater knowledge gains compared to those in the single-agent condition. This effect was observed regardless of participants' health literacy levels, contrasting prior work where an intervention with a single virtual agent only improved knowledge for participants with adequate health literacy [14]. Our work suggests that multi-agent interventions can be an effective strategy for enhancing learning for both inadequate and adequate health literacy users, highlighting the potential of multi-agent interventions to promote more accessible and equitable health education.

Despite results supporting a multi-agent intervention's overall knowledge improvement, our findings did not find any significant attitude improvement. Therefore, *H2* is not supported. Other research in health communication has found similar trends, and suggests this may be because attitudes are often deeper-seated than knowledge, and thus more difficult to change [2, 26]. Even when new information is introduced, people tend to retain existing beliefs, making surface-level interventions less effective at shifting attitudes [2]. Therefore, increasing the number of virtual agents in an educational health intervention may not be sufficient to change attitudes towards clinical trials.

Our results reject *H3*, indicating no significant intention improvement in the multi-agent condition compared to the single-agent condition. However, we observed a positive trend when accounting for users' health literacy. Specifically, the impact of the multi-agent intervention on intention improvement varied by health literacy level: participants with adequate health literacy showed a greater increase in intention scores compared to those with inadequate health literacy in the multi-agent condition. This trend may be linked to participants with adequate health literacy reporting significantly higher trust in the virtual agents under the multi-agent condition, as health communication research suggests that a "trusting relationship" with healthcare professionals can motivate cancer patients to participate in clinical trials [39]. However, more research is necessary to better understand the relationship between trust in virtual agents and users' intention to engage in health-related behaviors such as clinical trial participation.

In addition to multiple virtual agents' impact on knowledge, attitudes, and intention, our work also offers insights into users' preferences for virtual agents' information delivery. Notably, our study revealed that participants with inadequate health literacy significantly gravitated towards information being explained using less technical language. This observation was especially pronounced in the multi-agent condition. However, it remains unclear if the multi-agent condition or health literacy was responsible for this behavior. This observation warrants future investigation into whether a multi-agent intervention that models communicative support typically provided by real-world companions [11, 27, 45] can better facilitate information delivery to match users' needs.

6.1 Implications for Virtual Agent Developers

Our work contributes to intelligent virtual agent research by advancing understanding of how multi-agent interventions can address clinical trial barriers and how their impact varies for different kinds of users. Based on our findings, we recommend that developers of virtual agent interventions consider leveraging multiple agents when the goal is to broaden access to health information and support inclusive knowledge gains across health literacy levels. Additionally, preliminary trends suggest multi-agent interventions may boost behavioral intentions more than single-agent systems for users with higher health literacy. This may relate to greater trust in multi-agent approaches, but given the marginal significance and correlational nature of the data, these findings should be interpreted cautiously. Ultimately, our work suggests that multi-agent virtual interventions offer a promising approach to addressing cancer clinical trial participation, highlighting their potential use for other health contexts.

7 Limitations & Future Work

Our findings offer insights into the impact of multi-agent virtual systems on clinical trial-related knowledge, attitudes, and intentions, but several limitations remain. First, we did not isolate specific features of the multi-agent condition (e.g., agent role or gender), so it is unclear which aspects drove the observed effects. Nonetheless, comparing single- and multi-agent designs holistically provides valuable insight into their overall impact. Second, we relied on self-reported health literacy, which may be subject to bias. Our participant sample skewed toward higher health literacy and was predominantly White and female, limiting generalizability. Despite this, consistent patterns observed across groups still offer meaningful design implications. Third, we did not assess participants' attitudes toward technology or previous experience with virtual agents, which could have influenced responses; future work should account for these factors. Finally, the virtual agents' neural text-to-speech and limited nonverbal behavior may have reduced perceived naturalness, though these characteristics were kept consistent across conditions.

This work opens up several interesting avenues for future investigations. Examining the influence of the specific roles that virtual agents should take on within a multi-agent system, such as providing technical information, offering emotional support, or guiding users through complex decision-making, could help improve user outcomes and accessibility for low health literacy users. Longitudinal studies with multi-agent systems could also offer valuable insights into the long-term impact of such systems on users' clinical trial participation, assessing how sustained interaction with multiple agents influences attitudes, behaviors, and knowledge over time. Finally, extending multi-agent interventions to other contexts and populations, such as cancer screening and mental health, could expand our understanding of how these interventions can be leveraged in other domains. We call for future research to further explore how multi-agent systems can be refined and adapted for diverse users across a wider range of health applications.

8 Conclusion

Our work explored how multi-agent interventions can be used to address barriers to clinical trial participation and how their impact varies across users with different levels of health literacy. Our findings demonstrate that multi-agent virtual interventions are a promising strategy for improving knowledge outcomes and may also enhance behavioral intentions for individuals with higher health literacy. However, their effectiveness in shifting attitudes is limited, suggesting that simply increasing the number of agents may not be sufficient to influence deeper-seated beliefs. Ultimately, our work provides initial insights and potential opportunities for using multi-agent interventions to address barriers to cancer clinical trial participation, an important health communication challenge.

References

[1] Abeer Alam, Michael Lwin, Aila Khan, and Omar Mubin. 2024. Impact of Robot Size and Number on Human–Robot Persuasion. *Information* 15, 12 (2024), 782.

[2] Dolores Albarracin and Sharon Shavitt. 2018. Attitudes and attitude change. *Annual review of psychology* 69, 1 (2018), 299–327.

[3] Terrance L Albrecht, Susan S Eggly, Marci EJ Gleason, Felicity WK Harper, Tanina S Foster, Amy M Peterson, Heather Orom, Louis A Penner, and John C Ruckdeschel. 2008. Influence of clinical communication on patients' decision making on participation in clinical trials. *Journal of Clinical Oncology* 26, 16 (2008), 2666–2673.

[4] Allison J. Applebaum. 2019. *Cancer Caregivers.* Oxford University Press. doi:10.1093/med/9780190868567.001.0001

[5] Tessa Beinema, Harm op den Akker, Lex van Velsen, and Hermie Hermens. 2021. Tailoring coaching strategies to users' motivation in a multi-agent health coaching application. *Computers in Human Behavior* 121 (2021), 106787.

[6] Timothy W Bickmore, Laura M Pfeifer, Donna Byron, Shaula Forsythe, Lori E Henault, Brian W Jack, Rebecca Silliman, and Michael K Paasche-Orlow. 2010. Usability of conversational agents by patients with inadequate health literacy: evidence from two clinical trials. *Journal of health communication* 15, S2 (2010), 197–210.

[7] Timothy W Bickmore, Laura M Pfeifer, and Brian W Jack. 2009. Taking the time to care: empowering low health literacy hospital patients with virtual nurse agents. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 1265–1274.

[8] Timothy W Bickmore, Laura M Pfeifer, and Michael K Paasche-Orlow. 2009. Using computer agents to explain medical documents to patients with low health literacy. *Patient education and counseling* 75, 3 (2009), 315–320.

[9] Timothy W Bickmore, Dina Utami, Robin Matsuyama, and Michael K Paasche-Orlow. 2016. Improving access to online health information with conversational agents: a randomized controlled experiment. *Journal of medical Internet research* 18, 1 (2016), e1.

[10] Emma G Bryan, Huan Chen, Melissa Vilaro, Haoran Chu, Gabriella Grillo, Palani Te, Miriam Buhr, Stephen Anton, and Janice L Krieger. 2025. Developing a supportive virtual human to deliver clinical trial education for older women and other populations historically excluded from research. *Patient Education and Counseling* 130 (2025), 108485.

[11] Silvie Cooper and Fiona Stevenson. 2022. Communicating decisions about care with patients and companions in emergency department consultations. *Health Expectations* 25, 4 (2022), 1766–1775.

[12] AL Dias, JH Chao, D Lee, Y Wu, and GH Kloecker. 2016. Patient perceptions concerning clinical trials in oncology patients. *Contemporary Clinical Trials Communications* 4 (2016), 179–185.

[13] James N Dionne-Odom, Erin E Kent, Gabrielle B Rocque, Andres Azuero, Erin R Harrell, Shena Gazaway, Rhiannon D Reed, Reed W Bratches, Avery C Bechthold, Kyungmi Lee, et al. 2023. Family caregiver roles and challenges in assisting patients with cancer treatment decision-making: Analysis of data from a national survey. *Health Expectations* 26, 5 (2023), 1965–1976.

[14] Joy Egede, Maria J Galvez Trigo, Adrian Hazzard, Martin Porcheron, Edgar Bodiaj, Joel E Fischer, Chris Greenhalgh, and Michel Valstar. 2021. Designing an adaptive embodied conversational agent for health literacy: a user study. In *Proceedings of the 21st ACM International Conference on Intelligent Virtual Agents*. 112–119.

[15] Susan Eggly, Terrance L Albrecht, Felicity WK Harper, Tanina Foster, Melissa M Franks, and John C Ruckdeschel. 2008. Oncologists' recommendations of clinical trial participation to patients. *Patient Education and Counseling* 70, 1 (2008), 143–148.

[16] Osmane El Montaser, Sébastien Lallé, and Mathieu Muratet. 2025. Evaluating the Impact of Automated Hints in a 3D Educational Escape Game: A Comparative Study of Accessibility and Computer Science Versions. In *Proceedings of the 30th International Conference on Intelligent User Interfaces*. 685–694.

[17] Peter M Ellis, Phyllis N Butow, Martin HN Tattersall, Stewart M Dunn, and Nehmat Houssami. 2001. Randomized clinical trials in oncology: understanding and attitudes predict willingness to participate. *Journal of Clinical Oncology* 19, 15 (2001), 3554–3561.

[18] Mina Fallah, Timothy Bickmore, Stefan Olafsson, Michael Paasche-Orlow, Andrew Joseph Mrkva, and Jared W Magnani. 2024. Exploring the Potential of Virtual Agents in Atrial Fibrillation Management: Insights from a Randomized Trial. In *Proceedings of the 24th ACM International Conference on Intelligent Virtual Agents*. 1–9.

[19] Google Cloud. 2025. *Text-to-Speech API.* https://cloud.google.com/text-to-speech

[20] Jolie Haun, Stephen Luther, Virginia Dodd, and Patricia Donaldson. 2012. Measurement variation across health literacy assessments: implications for assessment selection in research and practice. *Journal of health communication* 17, sup3 (2012), 141–159.

[21] Jolie Haun, Virginia Noland-Dodd, Jill Varnes, John Graham-Pole, Barbara Rienzo, and Patricia Donaldson. 2009. Testing the BRIEF health literacy screening tool. *Fed Pract* 26, 12 (2009), 24–31.

[22] Chloe E Holden, Sally Wheelwright, Amélie Harle, and Richard Wagland. 2021. The role of health literacy in cancer care: a mixed studies systematic review. *PloS one* 16, 11 (2021), e0259815.

[23] Reshmashree B Kantharaju, Dominic De Franco, Alison Pease, and Catherine Pelachaud. 2018. Is two better than one? Effects of multiple agents on user persuasion. In *Proceedings of the 18th international conference on intelligent virtual agents*. 255–262.

[24] Reshmashree B Kantharaju, Alison Pease, Dennis Reidsma, Catherine Pelachaud, Mark Snaith, Merijn Bruijnes, Randy Klaassen, Tessa Beinema, Gerwin Huizing, Donatella Simonetti, et al. 2019. Integrating argumentation with social conversation between multiple virtual coaches. In *Proceedings of the 19th ACM International Conference on Intelligent Virtual Agents*. 203–205.

[25] Debbie Isobel Keeling, Angus Laing, and Ko De Ruyter. 2018. Evolving roles and structures of triadic engagement in healthcare. *Journal of Service Management* 29, 3 (2018), 352–377.

[26] Alexandra Y Kruse, Lise L Kjaergard, Kim Krogsgaard, Christian Gluud, Erik L Mortensen, Adam Gottschau, Anders M Bjerg, et al. 2000. A randomized trial assessing the impact of written information on outpatients' knowledge about and attitude toward randomized clinical trials. *Controlled Clinical Trials* 21, 3 (2000), 223–240.

[27] Rebekah C Laidsaar-Powell, Phyllis N Butow, Stella Bu, Cathy Charles, Amiram Gafni, Wendy WT Lam, Jesse Jansen, Kirsten J McCaffery, Heather L Shepherd, Martin HN Tattersall, et al. 2013. Physician–patient–companion communication and decision-making: a systematic review of triadic medical consultations. *Patient education and counseling* 91, 1 (2013), 3–13.

[28] Aisha T Langford, Sarah T Hawley, Sue Stableford, Jamie L Studts, and Margaret M Byrne. 2020. Development of a plain language decision support tool for cancer clinical trials: blending health literacy, academic research, and minority patient perspectives. *Journal of Cancer Education* 35 (2020), 454–461.

[29] Berit Hofset Larsen, Jennifer Gerwing, Tonje Lundeby, Pål Gulbrandsen, and Reidun Førde. 2025. More than an extra set of ears: A video-based analysis exploring the role of companions in routine oncology visits. *Journal of Psychosocial Oncology* (2025), 1–22.

[30] Guek Eng Lee, Mandy Ow, Desiree Lie, and Rebecca Dent. 2016. Barriers and facilitators for clinical trial participation among diverse Asian patients with breast cancer: a qualitative study. *BMC women's health* 16 (2016), 1–8.

[31] Shoou-Yih D Lee, Ahsan M Arozullah, and Young Ik Cho. 2004. Health literacy, social support, and health: a research agenda. *Social science & medicine* 58, 7

(2004), 1309–1321.
[32] Jessica Middlemis Maher, Jonathan C Markey, and Diane Ebert-May. 2013. The other half of the story: effect size analysis in quantitative research. *CBE—Life Sciences Education* 12, 3 (2013), 345–351.
[33] Julien Mancini, Jean Genève, Florence Dalenc, Dominique Genre, Alain Monnier, Pierre Kerbrat, Rémy Largillier, Daniel Serin, Maria Rios, Henri Roché, et al. 2007. Decision-making and breast cancer clinical trials: how experience challenges attitudes. *Contemporary clinical trials* 28, 6 (2007), 684–694.
[34] Andrew Maxim, Mohan Zalake, and Benjamin Lok. 2023. The Impact of Virtual Human Vocal Personality on Establishing Rapport: A Study on Promoting Mental Wellness Through Extroversion and Vocalics. In *Proceedings of the 23rd ACM International Conference on Intelligent Virtual Agents*. 1–8.
[35] Neal J Meropol, Yu-Ning Wong, Terrance Albrecht, Sharon Manne, Suzanne M Miller, Anne Lederman Flamm, Al Bowen Benson III, Joanne Buzaglo, Michael Collins, Brian Egleston, et al. 2016. Randomized trial of a web-based intervention to address barriers to clinical trials. *Journal of Clinical Oncology* 34, 5 (2016), 469–478.
[36] MET4Citizen. 2025. *TalkingHead: Real-Time 3D Avatars for Embodied Agents*. https://github.com/met4citizen/TalkingHead
[37] Sing Yu Moorcraft, Cheryl Marriott, Clare Peckitt, David Cunningham, Ian Chau, Naureen Starling, David Watkins, and Sheela Rao. 2016. Patients' willingness to participate in clinical trials and their views on aspects of cancer research: results of a prospective patient survey. *Trials* 17 (2016), 1–12.
[38] Sharon Mozgai, Arno Hartholt, Dayo Akinyemi, Katarina Kubicek, Albert Rizzo, and Michele Kipke. 2020. Development and Initial Feasibility Testing of the Virtual Research Navigator (VRN): A Public-Facing Agent-Based Educational System for Clinical Research Participation. In *HCI International 2020-Posters: 22nd International Conference, HCII 2020, Copenhagen, Denmark, July 19–24, 2020, Proceedings, Part II 22*. Springer, 304–307.
[39] Mary Murphy, Eilís McCaughan, Gareth Thompson, Matthew A Carson, Jeffrey R Hanna, Monica Donovan, Richard H Wilson, and Donna Fitzsimons. 2023. Trusting relationships between patients with non-curative cancer and healthcare professionals create ethical obstacles for informed consent in clinical trials: a grounded theory study. *BMC Palliative Care* 22, 1 (2023), 85.
[40] Clifford Nass, Jonathan Steuer, and Ellen R Tauber. 1994. Computers are social actors. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 72–78.
[41] National Cancer Institute. 2025. *Participate in Cancer Research*. https://www.cancer.gov/research/participate
[42] Farnaz Nouraei, Keith Rebello, Mina Fallah, Prasanth Murali, Haley Matuszak, Valerie Jap, Andrea Parker, Michael Paasche-Orlow, and Timothy Bickmore. 2025. Virtual Agent-Based Communication Skills Training to Facilitate Health Persuasion Among Peers. *Proceedings of the ACM on Human-Computer Interaction* 9, 2 (2025), 1–24.
[43] OpenAI. 2025. OpenAI Platform: Assistants Overview. https://platform.openai.com/docs/assistants/overview. Accessed: April 16, 2025.
[44] Maryke Peter, Stacy Maddocks, Clarice Tang, and Pat G Camp. 2024. Simplicity: using the power of plain language to encourage patient-centered communication. *Physical therapy* 104, 1 (2024), pzad103.
[45] Marco Pino and Victoria Land. 2022. How companions speak on patients' behalf without undermining their autonomy: Findings from a conversation analytic study of palliative care consultations. *Sociology of Health & Illness* 44, 2 (2022), 395–415.
[46] Plain Language Action and Information Network. 2025. *Federal Plain Language Guidelines*. https://www.plainlanguage.gov/guidelines/
[47] Mary C Politi, Marie D Kuzemchak, Kimberly A Kaphingst, Hannah Perkins, Jingxia Liu, and Margaret M Byrne. 2016. Decision aids can support cancer clinical trials decisions: results of a randomized trial. *The oncologist* 21, 12 (2016), 1461–1470.
[48] Prolific Academic Ltd. 2025. *Prolific*. https://www.prolific.com/
[49] Ready Player Me. 2025. *Ready Player Me*. https://readyplayer.me/
[50] Byron Reeves and Clifford Nass. 1996. The media equation: How people treat computers, television, and new media like real people. *Cambridge, UK* 10, 10 (1996), 19–36.
[51] David A Ross, Peter G Smith, and Richard H Morrow. 2015. *Chapter 2 Types of Intervention and Their Development*. Oxford University Press.
[52] Emre Sezgin. 2024. Redefining virtual assistants in health care: the future with large language models. e53225 pages.
[53] Ian Steenstra, Prasanth Murali, Rebecca Perkins, Natalie Joseph, Michael Paasche-Orlow, and Timothy Bickmore. 2023. Changing Parent Attitudes Towards HPV Vaccination by Including Adolescents in Multiparty Counseling using Virtual Agents. In *Proceedings of the 23rd ACM International Conference on Intelligent Virtual Agents*. 1–8.
[54] Ana Suarez Vazquez, Leticia Suárez Álvarez, and Ana Belen Del Rio Lanza. 2022. Communicating with companions. The impact of companion empowerment and companion literacy on the well-being of elderly patients. *Health Communication* 37, 5 (2022), 648–655.
[55] Thitaree Tanprasert, Sidney S Fels, Luanne Sinnamon, and Dongwook Yoon. 2024. Debate chatbots to facilitate critical thinking on youtube: Social identity and conversational style make a difference. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–24.
[56] Amelia A Trant, Lucas Walz, Whitney Allen, Jose DeJesus, Christos Hatzis, and Andrea Silber. 2020. Increasing accrual of minority patients in breast cancer clinical trials. *Breast Cancer Research and Treatment* 184 (2020), 499–505.
[57] Joseph M Unger, Dawn L Hershman, Cathee Till, Lori M Minasian, Raymond U Osarogiagbon, Mark E Fleury, and Riha Vaidya. 2021. "When offered to participate": a systematic review and meta-analysis of patient agreement to participate in cancer clinical trials. *JNCI: Journal of the National Cancer Institute* 113, 3 (2021), 244–257.
[58] Joseph M Unger, Lawrence N Shulman, Matthew A Facktor, Heidi Nelson, and Mark E Fleury. 2024. National estimates of the participation of patients with cancer in clinical research studies based on Commission on cancer accreditation data. *Journal of Clinical Oncology* (2024), JCO–23.
[59] Anna Winkler, Patrick Kutschar, Stefan Pitzer, Antje van der Zee-Neuen, Susanne Kerner, Jürgen Osterbrink, and Simon Krutter. 2023. Avatar and virtual agent-assisted telecare for patients in their homes: A scoping review. *Journal of Telemedicine and Telecare* (2023), 1357633X231174484.
[60] Jacob O Wobbrock, Leah Findlater, Darren Gergle, and James J Higgins. 2011. The aligned rank transform for nonparametric factorial analyses using only anova procedures. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 143–146.
[61] Mohan Zalake, Alexandre Gomes De Siqueira, Krishna Vaddiparti, Pavlo Antonenko, and Benjamin Lok. 2021. Towards understanding how virtual human's verbal persuasion strategies influence user intentions to perform health behavior. In *Proceedings of the 21st ACM International Conference on Intelligent Virtual Agents*. 216–223.
[62] Mohan Zalake, Fatemeh Tavassoli, Lauren Griffin, Janice Krieger, and Benjamin Lok. 2019. Internet-based tailored virtual human health intervention to promote colorectal cancer screening: design guidelines from two user studies. In *Proceedings of the 19th ACM International Conference on Intelligent Virtual Agents*. 73–80.
[63] Zhe Zhang and Timothy Bickmore. 2018. Medical shared decision making with a virtual agent. In *Proceedings of the 18th International Conference on Intelligent Virtual Agents*. 113–118.
[64] Shuo Zhou and Timothy Bickmore. 2022. A virtual counselor for breast cancer genetic counseling: adaptive pedagogy leads to greater knowledge gain. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–17.

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009