Empowering Diverse Health Literacy Needs with Multi-Agent Health Interventions: Design Challenges for Socially Aware Multi-Agent Systems

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

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

Abstract

Virtual agents are increasingly used in healthcare to deliver information and support decision-making. Yet most real-world clinical encounters involve multiple participants such as doctors, nurses, and caregivers. Multi-agent systems offer a promising paradigm for healthcare interventions, as they can better model these social dynamics. To explore their potential, we conducted a study comparing interactions with one versus two virtual agents, accounting for participants' health literacy. We examined attitudes toward decision making across three outcomes. Results suggest that multi-agent interventions may reduce decision-making pressure than single-agent interventions. However, users' health literacy plays a bigger role regarding information and satisfaction levels. Based on these preliminary results, we discuss challenges and opportunities for building socially aware multi-agent systems to empower users of diverse health literacy needs.

CCS Concepts

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

Keywords

virtual agents, health literacy, multi-agent systems

ACM Reference Format:
Rashi Ghosh and Benjamin Lok. 2025. Empowering Diverse Health Literacy Needs with Multi-Agent Health Interventions: Design Challenges for Socially Aware Multi-Agent Systems. In *Proceedings of the 13th International Conference on Human-Agent Interaction (HAI '25), November 10–13, 2025, Yokohama, Japan.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3765766.3765769

1 Introduction

Virtual agents—computer-generated characters that appear and behave like real humans [38]—are increasingly used in healthcare to promote healthy behaviors across diverse populations [33, 43, 44]. Through interactive, personalized conversations, they can enhance user understanding [6], build trust [40], and support health-related decision-making [42]. Despite these advances, most virtual health interventions rely on a single-agent paradigm, where one agent either adopts a narrow role (e.g., expert or peer) or attempts to embody multiple roles simultaneously. In contrast, healthcare communication in real-world contexts rarely unfolds with a single voice. Patients often interact with multiple social actors [14, 18, 19] including physicians, family members, and caregivers. Multi-agent systems provide a pathway to simulate these dynamics, potentially creating more realistic, engaging, and effective interventions.

One critical application is improving participation in cancer clinical trials, where adult enrollment remains persistently low (2–8%) [36, 37]. Barriers include limited awareness and poor communication, particularly for individuals with low health literacy who may struggle to interpret complex trial information [17, 22–24, 35]. Prior work has shown that single-agent interventions can effectively communicate complex health information across literacy levels [3, 45], including for low health literacy populations [4]. At the same time, single-agent systems have been reported as more difficult to interact with for lower health literacy users [11] while often providing greater comprehension benefits for higher health literacy users [5]. These findings suggest that single-agent interventions may not sufficiently support the diverse needs of users, particularly in contexts where decision-making is complex and socially mediated, such as clinical trial participation.

To address this, we explore the potential of multi-agent interventions that better reflect the collaborative and multi-voiced nature of real-world healthcare communication. In practice, patients are often accompanied by a third person companion during medical appointments, and such companions have been shown to improve comprehension and engagement by clarifying or contextualizing physicians' explanations [8, 19, 21, 30, 34]. In this paper, we present findings from a study comparing a web-based intervention with one virtual agent (doctor) versus two virtual agents (doctor and companion) across health literacy groups. Participants interacted with the virtual agents to learn about clinical trials, and reported their perceived pressure, amount of information received, and satisfaction as important metrics of decision-making support [44]. Our findings reveal that multi-agent systems have potential to reduce pressure during decision-making. However, users' health literacy plays a stronger role regarding information and satisfaction levels. Our findings highlight challenges and foster discussion for building socially aware multi-agent systems for healthcare interventions. We identify open design challenges for socially aware multi-agent interventions such as defining complementary roles (e.g., medical expert, empathetic peer) and orchestrating turn-taking strategies between multiple agents to best support users. By connecting our results to design considerations, we advance the discussion on how multi-agent interventions can contribute to the next generation of socially intelligent, empathetic health technologies.

HAI '25, Yokohama, Japan

ACM ISBN 979-8-4007-2178-6/25/11
https://doi.org/10.1145/3765766.3765769

2 Study Design

2.1 Virtual Agent System

We developed a web-based intervention with two conditions: a single-agent system featuring a virtual doctor (Dr. Alex) and a multi-agent system featuring a virtual oncologist and a virtual companion (Dr. Alex and Jordan). The doctor role represented an oncologist [1, 9, 12], who often introduces clinical trials, while the companion role reflected real-world caregivers who support patients and help process medical information [2, 10].

The virtual agents were presented as 3D characters, shown waist-up, with synthesized voices and basic nonverbal behaviors. 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 1), similar to other virtual agent studies for health-related interventions [25] [41].

All dialogue was pre-scripted and adapted from an established clinical trials intervention [26]. To address different information needs, response options were pre-generated at three levels of language complexity:

- *Base* - Information directly from the National Cancer Institute (NCI) [27] with no modification
- *Less Technical* - Modifications to the *base* response using simpler language and less medical jargon based on plain language guidelines [31], which is the recommended practice for communicating lower health literacy [22, 29]
- *More Technical* - Modifications to the *base* response using more complex language, medical jargon, and technical terms

The *less technical* and *more technical* modifications were generated using OpenAI's Assistants API [28], starting with the *base* response. All responses were validated by health communication experts.

At the start of the intervention, participants were guided through three questions (*What is a clinical trial*, *What is an IRB*, and *What is informed consent*? [26]). For each question, participants were shown the three explanation options with key language differences bolded and asked to select which option they preferred: base, less technical, or more technical. In the single-agent condition, Dr. Alex presented the explanation options and responded based on the user's selection. In the multi-agent condition, Jordan presented the explanation options, and Dr. Alex responded based on the user's selection.

Once familiar with the different explanation options, a preferred explanation type was determined. Participants proceeded to the question-answering phase where participants asked seven additional questions about clinical trials. Dr. Alex always responded to the user in the preferred explanation type. Participants could change the explanation style at any point during the intervention by clicking on Jordan in the multi-agent condition, and Dr. Alex in the single-agent condition.

2.2 Procedure

We conducted a between-participants study using Prolific, an online research recruitment platform [32] to examine the impact of a multi-agent intervention addressing barriers to clinical trial participation. Participants were randomly assigned to the single-agent condition (doctor) multi-agent condition (doctor and companion). Dialogue and content were held constant across conditions; only the number of agents varied. The study was approved by the university's IRB (IRB#ET00045896). A total of sixty-one U.S. adults over the age of 18 with a current or past cancer diagnosis were recruited. Participants first completed a pre-survey on Qualtrics, where they viewed the informed consent and completed a validated health literacy questionnaire [16], which was used to group them as low or high health literacy. Then, participants were redirected to the intervention website. Finally, participants completed the post-survey on Qualtrics, where they responded to three questions regarding decision-making outcomes [44]: perceived amount of information received, perceived pressure during the intervention, and satisfaction with the intervention. The entire study lasted about 25 minutes. All participants received monetary compensation of $5 USD.

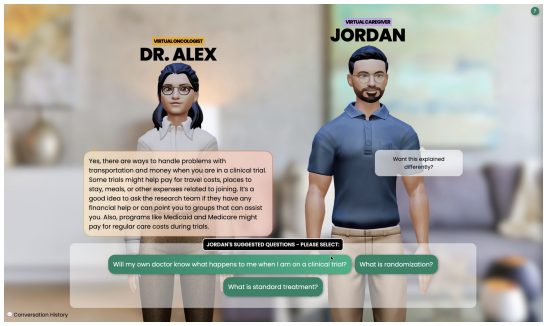

Figure 1: Multiple-choice button input interface during the question-asking phase for the multi-agent intervention

3 Results

Data analysis was conducted using R. Due to violation of normality assumptions assessed by Shaprio-Wilk tests, we used the Align Rank Transform (ART) ANOVA for nonparametric factorial analysis [39], similar to other work [13, 25]. We included Condition (single vs. multi-agent) and Health Literacy (low vs. high) as factors. Participants were grouped as either low or high health literacy based on their scores from the health literacy questionnaire, which is common practice in health communication literature [15] [5].

Perceived Pressure. Condition had a main effect on how much pressure participants felt during the intervention, where those in the multi-agent condition reported feeling significantly less pressured compared to those in the single-agent condition ($F(1, 57) = 4.36, p = .041, \eta^2 = .07$). Health literacy also had a main effect, where participants with low health literacy reported significanlty less pressure than those with high health literacy ($F(1, 57) = 16.99, p < .001, \eta^2 = .23$).

Perceived Amount of Information. Condition had no main effect on how much information participants felt they received during the intervention. Health literacy had a main effect, where

participants with low health literacy felt they received more information than those with high health literacy ($F(1, 57) = 4.71, p = .034, \eta^2 = .08$).

Satisfaction with Decision-Making. Condition had no main effect satisfaction. Health literacy had a main effect, where participants with high health literacy reported greater satisfaction with the decision-making process compared to those with low health literacy, ($F(1, 57) = 5.03, p = .029, \eta^2 = .08$).

4 Discussion

Our preliminary findings suggest that multi-agent interventions can reduce feelings of pressure during health-related decision-making, aligning with prior evidence that companions in real-world clinical encounters help patients feel more supported. By contrast, perceptions of the amount of information received and satisfaction with the decision-making process appeared to be shaped more strongly by users' health literacy than by the number of agents. We discuss our findings in this section, hilighting persistent challenges in designing inclusive decision-support systems and opportunities for incorporating social awareness into the design of multi-agent systems.

Interestingly, although participants with low health literacy felt they received more information compared to those with high health literacy, they also reported lower satisfaction. This finding aligns with prior work where users with lower health literacy have been shown to have more difficulty interacting with adaptive virtual agents [11], while users with higher health literacy benefited more in terms of comprehension with virtual agents [5]. Our results may highlight low health literacy users experiencing information overload. The Cognitive Theory of Multimedia Learning warns that more complex interfaces, such as virtual agent interventions, might be too distracting [7], and healthcare literature identifies low health literacy as a predictor of health information overload [20]. Therefore, it is possible that low health literacy participants perceiving the increased amount of information in the virtual agent intervention, regardless of number of agents, was due to cognitive overload.

Future work could refine the role of a second virtual agent to detect if users are feeling overwhelmed or confused by the amount of information they are receiving, and intervene accordingly. Our intervention was restricted to multiple choice input, but future interventions could allow for free-text or voice input to enable methods such natural language processing and analysis of vocalic properties to detect users' emotional and cognitive state. This could help the system infer confusion or disengagement, triggering timely intervention from the second virtual agent. For example, for lower literacy users, the companion agent might take a more proactive role, offering scaffolding, prompts, or clarifications, while for higher literacy users it could remain reactive, allowing greater autonomy. These strategies would make multi-agent interventions more socially aware and able to adapt to diverse health literacy needs, empowering more users.

5 Conclusion

This work provides preliminary evidence that multi-agent virtual health interventions can shape decision-making experiences in meaningful ways. While adding a companion agent reduced perceived pressure, outcomes such as information and satisfaction levels varied more strongly by health literacy. These findings point to the potential of multi-agent systems to create more socially attuned health technologies, but also reveal gaps in addressing diverse user needs. Based on our findings, we highlight next steps for making multi-agent systems more socially aware by detecting and responding to user cues that might signal confusion and mental overload. This could enable agents to better sense user states and tailor support in real time. By advancing these directions, we aim to inform the development of socially aware, cooperative multi-agent systems that improve accessibility in healthcare interventions for diverse users.

References

[1] 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.

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

[3] 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.

[4] 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.

[5] 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.

[6] 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.

[7] Juan C Castro-Alonso, Rachel M Wong, Olusola O Adesope, and Fred Paas. 2021. Effectiveness of multimedia pedagogical agents predicted by diverse theories: A meta-analysis. *Educational Psychology Review* 33, 3 (2021), 989–1015.

[8] 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.

[9] 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.

[10] 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.

[11] 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.

[12] 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.

[13] 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.

[14] Amy Fioramonte and Camilla Vásquez. 2019. Multi-party talk in the medical encounter: Socio-pragmatic functions of family members' contributions in the treatment advice phase. *Journal of Pragmatics* 139 (2019), 132–145.

[15] 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.

[16] 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.

[17] 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.

[18] Dominic Hurst and Sharon Mickan. 2017. Describing knowledge encounters in healthcare: a mixed studies systematic review and development of a classification. *Implementation Science* 12, 1 (2017), 35.

[19] 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.

[20] Israa Khaleel, Barbara C Wimmer, Gregory M Peterson, Syed Tabish Razi Zaidi, Erin Roehrer, Elizabeth Cummings, and Kenneth Lee. 2020. Health information overload among health consumers: A scoping review. *Patient education and counseling* 103, 1 (2020), 15–32.

[21] 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.

[22] 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.

[23] 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.

[24] 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.

[25] 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.

[26] 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.

[27] National Cancer Institute. 2025. *Participate in Cancer Research.* https://www.cancer.gov/research/participate

[28] OpenAI. 2025. OpenAI Platform: Assistants Overview. https://platform.openai.com/docs/assistants/overview. Accessed: January 20, 2025.

[29] 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.

[30] 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.

[31] Plain Language Action and Information Network. 2025. *Federal Plain Language Guidelines.* https://www.plainlanguage.gov/guidelines/

[32] Prolific Academic Ltd. 2025. *Prolific.* https://www.prolific.com/

[33] 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.

[34] 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.

[35] 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.

[36] 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.

[37] 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.

[38] 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.

[39] 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.

[40] Christopher You, Rashi Ghosh, Andrew Maxim, Jacob Stuart, Eric Cooks, and Benjamin Lok. 2022. How does a virtual human earn your trust? guidelines to improve willingness to self-disclose to intelligent virtual agents. In *Proceedings of the 22nd ACM International Conference on Intelligent Virtual Agents*. 1–8.

[41] 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.

[42] Mohan Zalake, Alexandre Gomes De Siqueira, Krishna Vaddiparti, Pavlo Antonenko, and Benjamin Lok. 2024. Evaluating the Effect of Adapting Virtual Humans Based on Individual Differences in Users. In *International Conference on Human-Computer Interaction*. Springer, 405–423.

[43] 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.

[44] 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.

[45] 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