Wizard of Oz with Strong Magnets: Exploring Haptic Interactions with Non-Humanoid AI Agents

Ahmet Börütecene

Linköping University, Media and Information Technology, SE-581 83, Linköping, Sweden

Abstract

Haptic interaction with artificial intelligence (AI) entities that have a physical embodiment is a developing area of research which investigates how to design communicative tangible expressions for human-AI interactions. However, traditional sketching techniques such as storyboards or making fully functional prototypes do not seem to be convenient options to capture these agents' potential for haptic interaction in context and rapidly explore possible ways of expression. In a recent study, we designed and built a Wizard of Oz (WOz) rig where users engaged in haptic interaction with a handheld non-humanoid robot as a form of embodied AI agent that was not working yet physically controlled with strong magnets by a human wizard hidden under the table. The paper reports on this WOz rig, shares reflections on the human wizard's bodily engagement with users through his magnet-driven magic wand, and discusses its potential as a low-tech responsive experimentation environment for exploring movement-based AI expressions. Our contribution is the WOz mechanism and design-related insights derived from its operation, which could be useful for similar studies on haptic interaction with embodied AI agents.

Keywords

Wizard of Oz, haptic interaction, embodied interaction, robotic artefacts, artificial intelligence, human-robot interaction

Introduction

Artificial intelligence (AI) is present in different forms in our everyday lives. We interact with AI entities mostly through visual (e.g., dashboards, editors) and auditory interfaces (e.g., voice assistants). Recently, exploration of haptic interactions with embodied AI agents emerged as a research direction in HCI [4, 21, 12] which explores how to design communicative tangible expressions for human-AI interactions. A well-known and broad category of embodied AI agents is robots.

However, the word robot has many mental images and ideals associated with it. Perhaps the most known and recurrent one represents an autonomous entity that looks and behaves exactly like a human, to such an extent that it is indistinguishable from an authentic one. Whether this ideal is worth pursuing or not, we should accept that today's technology does not enable us to design such a robot or embodied AI agent. We can also argue that interactions with humanoid agents are almost always delusional; the more an autonomous agent looks like a human, the higher the social expectation is. However, humanoid agents cannot afford such a sophisticated level of behavior promised by their appearance, at least not yet, whereas non-anthropomorphic robotic artefacts that demonstrate an elevated level of behavioral realism hold the potential for more coherent and engaging interactions [9].

In addition, human-AI interactions are mostly based on two sensory channels: visual/textual and auditory. However, humans also use their bodies and the haptic channel (sense of touch) to engage in direct interaction with the environment and objects of interest. Together with visual and auditory channels, the haptic channel allows us to form a nuanced understanding of interactions and situations. The sense of touch could support vision by providing another layer of information, and the information

Proceedings of ETIS 2022, November 7–10, 2022, Toulouse, France
EMAIL: ahmet.borutecene@liu.se
ORCID: 0000-0003-1153-1491

delivered from the haptic channel could reduce the load on vision [13, p.159]. Then, an alternative approach is to go in the opposite direction and design tangible, non-anthropomorphic AI agents that use nonverbal interactions such as expressive movements (i.e., movements performed to convey different states, emotions, intentions) to ground communication. Although these agents do not have a visual or auditory resemblance to a living form, humans attribute animacy and agency to them [14, 5]. This enables the agents to be perceived as autonomous entities that have their own behaviors, attitudes, and nonverbal language. These kinds of embodied AI agents mark an uncharted territory without established interaction elements, and the exploration of such non-idiomatic interactions requires a responsive experimentation environment to rapidly try out new behaviors and situations [11].

So, we asked a few questions, in close connection to the constellation of these research works and approaches mentioned, regarding the design of a physical AI agent that humans can literally get in touch with: What happens if humans interact with an embodied non-humanoid AI agent by placing their hand on its body to move it around? How would the agent communicate intent through the haptic channel? When the agent moves against their hand how would humans feel and interpret this physical movement? What about communicating emotions? These questions led us to adopt a commercial handheld robot with wheels as a non-humanoid AI agent and to construct a Wizard of Oz (WOz) rig to search for answers. In this regard, we considered that building a new mobile robot or employing an existing one with a form factor enabling manipulative gestures we were looking for and enough motor/wheel power for desired movements (such as showing resistance to the force applied by users) in haptic modality would have been costly and time consuming. Instead, we opted for a mechanism that would enable us to puppet an embodied AI and get a feel of how haptic interaction between humans and movement-based agents would play out [19]. We also imagined that such a rig could be used with non-digital prototypes (e.g., robot made of cardboard) to rapidly explore haptic human-AI interactions in terms of embodiment, expressions, and behavior. The WOz rig was thus based on the idea of using simple materials that would enable users to engage in haptic interaction with the robot which was not actually working yet physically controlled with strong magnets by a human wizard hidden under the table [20, 18].

In this paper, we first provide a brief account of the user study and describe the WOz rig. Then, we share reflections on the human wizard's bodily engagement with users through his magnet-driven magic wand and discuss the WOz rig's potential as a low-tech responsive experimentation environment for exploring movement-based AI expressions. In this report, our focus is on the WOz setup and its use, so we do not cover the details of the user study as it is beyond the scope of this paper.

Brief account of the user study

We designed a study to investigate how direct physical interaction with embodied AI agents influences the sense of agency (SoA) in humans in comparison to interacting with them via GUI. Our intention was to understand in which ways haptic modality contributes to the feeling of being in control during interaction with (semi-)autonomous agents in a social and collaborative context. The study comprised a game played over a fictional map on a table in which a robot called Rover and a user were teaming up to complete a mission on a foreign planet.

We ran 11 sessions (including three pilots), and in each session one participant and two researchers were involved. One researcher was responsible for managing the participants, conducting the game part, and running the retrospective think-aloud part and the other researcher was the author, responsible for playing the human wizard who was under the table during the study sessions.

The participants played the game standing in front of the table. Each participant was asked to complete two tasks together with the robot in two different modalities. One task was to collect the most valuable mineral among the three and the other was to reach the volcano on the map through a field of geysers without falling in them and collect a lava stone. In the haptic modality, the participants were supposed to place their hand on the robot resting on the map and physically move it by pushing, pulling, or rotating in a way to complete the tasks (e.g., reaching a destination, picking up an object etc.) In the GUI modality, they were able to move the robot in the same way, this time through a tablet screen by using a dedicated visual interface. In both modalities, the participants were starting the tasks from the base represented by a shield icon in the middle of the map (Figure 1), and they were moving towards items or places relevant for the task in question. A mineral was considered collected, or a destination reached, once the robot was positioned on them.

Figure **1**: (From left to right) The WOz rig with the robot resting on the table surface; controller lying in the space under the table where the wizard performs; close up on the robot resting on an early version of the game map. The shield icon in the center is the base where the game starts. The green item that lies behind the border made of cardboard is one of the valuable minerals available to collect.

While the participants were interacting using these modalities, the robot was expressing itself nonverbally through the ways it was moving, demonstrating agentive qualities. For example, if the robot did not agree with the direction its partner was moving in it was resisting against the pushing action performed on it (e.g., to communicate that the participant was moving towards the less valuable mineral). Or, if the robot wanted to hint at another object or path, it was taking its partner's hand towards that option instead of the other by moving its body (e.g., to communicate that the participant should move in another direction to go around a geyser).

After playing in each modality, the participants were filling in a seven-point Likert scale questionnaire that consisted of nine statements used to measure their SoA during the game (Figure 2). We created the questionnaire by adapting two different works—the items used by Endo et al. [3] and the scale created by Polito et al. [15]—to elicit feedback on how much control and engagement the participants felt in relation to the AI agent and the game. Once the participants completed the game, a retrospective think-aloud session was held. In this session, they were watching the videos of their interaction with the robot in each task which were recorded by a webcam from above the WOz rig. Finally, the researcher, as the wizard in charge of controlling the robot, was coming out of his hiding spot under the table for debriefing and a short, informal chat about the participant's experience. The collected data and the session videos were analyzed through the lens of the tool-agent spectrum by Rozendaal et al. [16] and the Social Agency continuum by Silver et al. [17] to understand how different modalities influenced the participants' SoA and their perception of the robot during interaction.

1. Rover moved just like I wanted it to, as if it were obeying my will.
2. I felt as if I were controlling the movement of Rover.
3. I felt as if I were causing the movement I saw.
4. Whenever I moved the sliders of the GUI, I expected Rover to move in the same way.
5. My experiences and actions felt self generated.
6. I embraced the suggestions [by the Rover] freely.
7. Following suggestions was hard.
8. I was mostly absorbed in what was going on.
9. I was reluctant to follow suggestions [from the Rover]

Figure **2**: The seven-point Likert scale (strongly disagree-strongly agree) questionnaire statements adapted from Endo et al. and Polito et al. Rover was the name of the robot in the game.

The preliminary results from the study suggest that the haptic modality could allow for high level of control over an embodied AI, fine-grained interactions, and enable users to get more detailed information about situations, although users may consider the latter sometimes excessive or confusing. On the other hand, interaction via GUI seems to provide less strong connection with the perceived agency of an embodied AI, although it provides sufficient control for interaction and may be considered suitable for cases where users do not need or want detailed information. The study enabled us to have an understanding of the movements' communicative, agentive, and expressive qualities, while the covert nature of the WOz setup helped us capture some valuable insights to consider when designing for haptic human-AI interactions such as the ghost agency effect (i.e., at some point one participant was thinking that he was being guided by the agent whereas he was the person who was steering it) and slightly unwelcoming attitude towards haptic embodied agents (i.e., one participant stated that she would have behaved more gently if she had known there was someone controlling the robot under the table).

In both modalities, the author was the wizard under the table who was performing these actions on the robot's behalf. Now let us see in more detail how the WOz rig was built, and how the wizard used it.

Magic Wand with Strong Magnets

For the study, we used a mobile robot called Root produced by iRobot for children to learn coding which had built-in magnets underneath [8]. We made a wooden hand controller with four strong magnets glued to its top part to leverage on the existing magnets inside the robot (Figure 3). These magnets were placed on the controller in such a way that when they came close to the table and were in line with the robot's magnets, they were strongly adhering to each other. In this way, the wizard could steer the robot while hiding under the table without touching it directly.

Figure **3**: (From left to right) View of the GUI and the robot wearing a crochet cover as used during the user study; view of the controller with strong magnets and the inverted map placed under the table; small cube magnet on the inverted map enabling the wizard to understand where the robot is.

However, the thickness of traditional desks or tables did not allow for the strong magnets to sufficiently adhere to the robot, so we made a special table building a wooden frame that was holding a thin plank on it which worked well. During the user study, the wizard was sitting on a mat under this special table we constructed which was held by two desks on each side. To prevent participants from thinking that there might be someone or a mechanism moving the robot from under the table, we used clothes to cover the WOz rig, trying to give an impression of a booth dedicated to the game (Figure 1).

However, we needed to develop the strong magnet controller further to adapt it to the study conditions. For example, we attached silicone rubber tape to each magnet on the hand controller to increase the adhesive force in sheer direction so that the steering action was easier and more controllable. Moreover, the controller had to move silently across the bottom surface of the table to prevent the study participants from thinking that someone might be steering the agent. As the plank was very thin, we needed to use a buffer material to mitigate the impact when the magnets adhered to the robot from under the plank. We placed kitchen sponges on top of the magnets which acted like a pillow that absorbed the first impact as well as provided the silence required during the interaction. We also used a thin cotton fabric to cover the sponges to reduce friction. The robot had a couple of visual cues such as eyes (displaying animation) and a nose-like button suggesting a kind of animal. As we wanted the users to engage in haptic interaction with the robot, we tweaked its appearance by using a crochet cover to hide these cues as these might have prevented them from focusing on its movements (Figure 3).

When the wizard was under the table, he needed to follow what was going on above him, so he had two laptops. One was connected to the webcam placed above the rig that was providing a video feed of the map and the participants. In this way, he was able to see where the robot was moving and act accordingly. On the surface under the table, there was an inverted version of the map lying on the table so that he knew approximately where the robot was at any time (Figure 3). In the haptic modality, he was using the magnet controller. When he saw in the video feed that the participant was steering the robot to a spot in the map where he should intervene, he was slowly and carefully approaching the controller to the plank surface to smoothly adhere to the robot. Once the controller was adhered, he could move, rotate, or let go of the robot (taking down the controller) according to the emerging situation for different purposes (e.g., showing resistance against the path chosen by the participant).

The other laptop was used to run the software that was establishing the connection between the robot and the tablet interface, which was enabling participants to use the browser-based visual interface to move the robot. In the GUI modality, the wizard was using his own smartphone to connect to the same dedicated website for the interface. In this way, he was also able to move the robot (without participants seeing any change on the interface) according to the information he was receiving from the video feed. So, when he was moving a slider in the interface, his command was overriding the participants' and vice versa. In addition to the video feed, he was placing a small magnet on the inverted map—on the spot that was matching the starting position of the robot in the GUI modality. This allowed him to have a more accurate view of where the robot was at any time on the map, by following the small magnet on the inverted map, and less divided attention while interacting with the visual interface. According to the game scenario, the robot-human team was supposed to go across a terrain on the map that contained dangerous areas to avoid. The participants could not see these areas on the surface map, but they were marked on the inverted map under the table. The wizard was using these circular marks to decide when (e.g., seeing that the small magnet approaching one of them) and where to maneuver the robot both in the haptic and GUI modality to help the participants navigate.

Being in the Magnetic Feel

This section is a retrospective account of what the researcher in charge of the robot experienced during the study as a wizard. First, it obviously felt a little bit uncomfortable for him staying under the table for almost an hour, trying to be attentive to several elements at the same time. Although he could see the light through the fabric of the tablecloth covering the WOz rig, it was feeling stuffy inside and he needed to be careful not to make any noise. Also, he had to make sure that he was not making any mistakes in maneuvering the robot, which required him to be very careful about where the robot was on the map. The thought of getting caught by the participants made him nervous from time to time, as he was physically so close to them and able to see their legs and feet approaching him through the fabric. In the end, out of 11 participants only one understood (as the debriefing revealed) that there was someone under the table controlling the robot. In the end of the study, some participants were not expecting to see him at all and were very much surprised that there was someone operating the robot under the table.

Although this was a bodily and cognitively tiresome process, he enjoyed the experience, as it enabled him to bodily perform as a physical AI agent. This slightly demanding WOz setup gave the researcher the possibility to get closer to the AI and channel his intentions through its embodiment while interacting with human participants. This made it possible to experiment with movement-based haptic AI expressions and see how they are perceived by humans. The sophisticated material and spatial setup added to the experience, as it created an engaging physical and emotional connection to the embodied agent, maybe more than only a GUI-based WOz setup would do.

Regarding bodily engagement, the wizard had some challenges. Sometimes, he could turn his body to a certain extent when he wanted to perform a move with the controller because he was in a restricted space. This made him think of the limited capacity of an AI entity and how it would struggle, bodily or cognitively, to interact with someone in collaborative situations. Dörrenbächer et al. [2] conducted a study where participants roleplayed robots, however in the study they did see each other and were aware of the roles. The fact that in our study the participants did not know that they were interacting with a human wizard under the table, and that the wizard should not make any noise to break the illusion, maybe created a more reserved, isolated space and thus an occasion for self-reflection for the person playing the wizard.

Finally, the WOz rig was a tight coupling of human body and physical material with some digital elements. In this setup, the wizard was literally in the loop as a human in haptic interactions between an AI agent and the users. This analogy may be more appropriate in exploratory scenarios with loose scripts where the wizard can iterate embodied agent's movements each time the scenario is played to rapidly create and observe new situations and outcomes. In this sense, this kind of WOz rig offers a low-tech (strong magnets, simple physical materials) yet responsive experimentation environment to relate to hi-tech artefacts and can be used as a mechanism to sensitize people from different disciplines to material, behavioral and aesthetic interaction qualities of embodied AI [10].

Discussion and conclusion

To summarize, we adopted a commercial robot in a WOz rig where it was physically manipulated in real-time by a human wizard—hidden under the table—to understand in which ways haptic modality contributes to the feeling of being in control during interaction with (semi-)autonomous agents as well as how users approach an embodied non-humanoid AI agent's expressive movements.

The WOz rig created a situation where the wizard was engaging in continuous bodily interaction with users thanks to the custom-made controller with strong magnets, although there was no direct touch between the two interactants. Physical feedback caused by the movements performed by users' hands on the robot such as pushing, pulling, and rotating enabled the wizard to have a haptic connection to the users as well as to the non-humanoid robot he was enacting. This was perfectly in line with the study's aim of exploring haptic interaction with embodied AI agents. The fact that this kind of WOz rig allows the interaction between the agent and the user to be in real-time, it has the potential to be configured for semi-structured, contextual scripts to sketch and prototype continuous human-AI interaction [1]. The magnet controller was not only mediating the manual actions performed on the robot above the wizard, but also conveying what the robot would feel in these moments of physical interaction and enabling the wizard to think through the artefact. The setup has also an influence on the wizard's awareness of his/her own body, its abilities and limits, flexibility and responsiveness as well as its relation to other bodies, which opens an interesting space for exploring felt experiences of joint movement and touch [6]. As we saw in the study, such a setup offers the touch and tangibility not only as a medium to design for but also as a medium to design with and through, by proposing a responsive experimentation environment where both users and wizards can experience multiple dimensions of tangible interaction (tangible manipulation, spatial interaction, embodied facilitation, expressive representation) [7] in real-time.

As future work, we plan to use the WOz rig this time in an overt setup where participants will be informed about the wizard's presence right from the beginning. Performing as a wizard, the researcher will be doing open-ended and transparent collective sketching with designers and domain experts to explore AI entities' agentive qualities and expressive movements. In connection to this, one potential of the WOz rig is configuring the robot to record a log of emerging interactions during these exploratory sketching sessions and use them as a resource to inform the design and implementation of expressive movements for an embodied agent, which could be used for further testing in real-life situations.

Acknowledgements

I would like to thank Marcus Jessen for planning and conducting the user study as part of his studies in the master's program in cognitive science at Linköping University and Mikael Pettersson for implementing the visual interface for the GUI to move the robot. Special thanks to our colleague Carl Nordlund for making the crochet cover for the robot.

References

[1] Niels van Berkel, Mikael B. Skov, and Jesper Kjeldskov. 2021. Human-AI interaction: intermittent, continuous, and proactive. Interactions 28, 6: 67–71. https://doi.org/10.1145/3486941

[2] Judith Dörrenbächer, Diana Löffler, and Marc Hassenzahl. 2020. Becoming a Robot - Overcoming Anthropomorphism with Techno-Mimesis. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20), 1–12. https://doi.org/10.1145/3313831.3376507

[3] Satoshi Endo, Jakob Fröhner, Selma Musić, Sandra Hirche, and Philipp Beckerle. 2020. Effect of External Force on Agency in Physical Human-Machine Interaction. Frontiers in Human

Neuroscience 14. Retrieved July 26, 2022 from https://www.frontiersin.org/articles/10.3389/fnhum.2020.00114
[4] Maliheh Ghajargar, Jeffrey Bardzell, Alison Marie Smith-Renner, Kristina Höök, and Peter Gall Krogh. 2022. Graspable AI: Physical Forms as Explanation Modality for Explainable AI. In Sixteenth International Conference on Tangible, Embedded, and Embodied Interaction (TEI '22). Association for Computing Machinery, New York, NY, USA, Article 53, 1–4. https://doi-org.e.bibl.liu.se/10.1145/3490149.3503666
[5] Guy Hoffman and Wendy Ju. 2014. Designing robots with movement in mind. Journal of Human-Robot Interaction 3, 1: 91–122. https://doi.org/10.5898/JHRI.3.1.Hoffman
[6] Kristina Höök, Caroline Hummels, Katherine Isbister, Patrizia Marti, Elena Márquez Segura, Martin Jonsson, Florian “Floyd” Mueller, Pedro A.N. Sanches, Thecla Schiphorst, Anna Ståhl, Dag Svanaes, Ambra Trotto, Marianne Graves Petersen, and Youn-kyung Lim. 2017. Soma-Based Design Theory. In Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA ’17), 550–557. https://doi.org/10.1145/3027063.3027082
[7] Eva Hornecker and Jacob Buur. 2006. Getting a grip on tangible interaction: a framework on physical space and social interaction. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI ’06), 437–446. https://doi.org/10.1145/1124772.1124838
[8] iRobot® Root®. Retrieved June 12, 2022 from https://www.irobot.com/en_US/irobot-root/Rt001.html
[9] Florent Levillain and Elisabetta Zibetti. 2017. Behavioral objects: the rise of the evocative machines. Journal of Human-Robot Interaction 6, 1: 4–24. https://doi.org/10.5898/JHRI.6.1.Levillain
[10] Jonas Löwgren. 2009. Toward an articulation of interaction esthetics. New Review of Hypermedia and Multimedia 15, 2: 129–146. https://doi.org/10.1080/13614560903117822
[11] Jonas Löwgren. 2016. On the significance of making in interaction design research. Interactions 23, 3: 26–33. https://doi.org/10.1145/2904376
[12] Maria Luce Lupetti, Cristina Zaga, and Nazli Cila. 2021. Designerly Ways of Knowing in HRI: Broadening the Scope of Design-oriented HRI Through the Concept of Intermediate-level Knowledge. In Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction (HRI ’21), 389–398. https://doi.org/10.1145/3434073.3444668
[13] Karon Maclean. 2008. Haptic Interaction Design for Everyday Interfaces. Reviews of Human Factors and Ergonomics 4: 149–194. https://doi.org/10.1518/155723408X342826
[14] Betti Marenko and Philip van Allen. 2016. Animistic design: how to reimagine digital interaction between the human and the nonhuman. Digital Creativity 27, 1: 52–70. https://doi.org/10.1080/14626268.2016.1145127
[15] Vince Polito, Amanda J. Barnier, and Erik Z. Woody. 2013. Developing the Sense of Agency Rating Scale (SOARS): An empirical measure of agency disruption in hypnosis. Consciousness and Cognition 22, 3: 684–696. https://doi.org/10.1016/j.concog.2013.04.003
[16] Marco C. Rozendaal, Evert van Beek, Pim Haselager, David Abbink, and Catholijn M. Jonker. 2020. Shift and Blend: Understanding the hybrid character of computing artefacts on a tool-agent spectrum. In Proceedings of the 8th International Conference on Human-Agent Interaction (HAI ’20), 171–178. https://doi.org/10.1145/3406499.3415066
[17] Crystal A. Silver, Benjamin W. Tatler, Ramakrishna Chakravarthi, and Bert Timmermans. 2020. Social Agency as a continuum. Psychonomic Bulletin & Review. https://doi.org/10.3758/s13423-020-01845-1
[18] Superflux. 2009. Sketch-a-Move. Retrieved June 12, 2022 from https://vimeo.com/5125096
[19] Peter Wang, Srinath Sibi, Brian Mok, and Wendy Ju. 2017. Marionette: Enabling On-Road Wizard-of-Oz Autonomous Driving Studies. In Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction (HRI ’17), 234–243. https://doi.org/10.1145/2909824.3020256
[20] Mechanical Turk. Wikipedia. Retrieved June 12, 2022 from https://en.wikipedia.org/w/index.php?title=Mechanical_Turk&oldid=1087252650
[21] Oren Zuckerman, Viva Sarah Press, Ehud Barda, Benny Megidish, and Hadas Erel. 2022. Tangible Collaboration: A Human-Centered Approach for Sharing Control With an Actuated-Interface. In

CHI Conference on Human Factors in Computing Systems (CHI '22), 1–13. https://doi.org/10.1145/3491102.3517449