
Spontaneous Interactions with a Virtually Embodied Intelligent Assistant in Minecraft

Fraser Allison¹

Microsoft Centre for Social NUI
University of Melbourne, Australia
fraser.allison@unimelb.edu.au

Ewa Luger¹

University of Edinburgh
Edinburgh, UK
eluger@exseed.ed.ac.uk

Katja Hofmann

Microsoft Research
Cambridge, UK
katja.hofmann@microsoft.com

¹Most of this research conducted while at Microsoft Research Cambridge, UK.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.
Copyright is held by the owner/author(s).
CHI'17 Extended Abstracts, May 06-11, 2017, Denver, CO, USA
ACM 978-1-4503-4656-6/17/05.
<http://dx.doi.org/10.1145/3027063.3053266>

Abstract

An increasing number of our technological interactions are mediated through virtually embodied characters and software agents powered by machine learning. To understand how users relate to and evaluate these types of interfaces, we designed a Wizard of Oz prototype of an embodied agent in Minecraft that learns from users' actions, and conducted a user study with 18 school-aged Minecraft players. We categorised nine main ways users spontaneously attempted to interact with and teach the agent: four using game controls, and five using natural language text input. This study lays groundwork for a better understanding of human interaction with learning agents in virtual worlds.

Author Keywords

Human-agent interaction; Minecraft; natural language.

ACM Classification Keywords

H.1.2. User/Machine Systems: Human factors.

Introduction

Semi-autonomous agents are an increasingly common element in our interactions with technology. These agents, or bots, are software programs that undertake tasks with little or no direct supervision by a user,

reacting somewhat independently to their context of operation. In interfaces with humans, they often take the form of virtual characters, whether as intelligent assistants such as Siri (Apple, 2011), natural language bots such as Xiaoice (Microsoft, 2014), or virtually embodied characters such as those found in digital games. The agent paradigm presents a different set of interaction design challenges than the previously dominant paradigm of direct manipulation of graphical user interfaces [14,18], as engaging with an autonomous or semi-autonomous agent requires a more collaborative working style and the development of trust in the agent's abilities [3,5,16].

An additional complication is added by the fact that so much recent progress in agent-based applications is founded on machine learning techniques. While machine learning has been stunningly successful in transforming the ability of computers to undertake a range of tasks, especially when generalisation or interpolation is needed, it also introduces risks and unpredictability into the human-machine relationship. As the behavioural rules that are shaped by machine learning are typically stochastic, its outcomes often cannot be guaranteed, but are instead only statistically predictable. In addition, although machine learning can be used to make an intelligent agent adapt to the user, there are questions about how to make this learning process transparent and intuitive to the user [8].

In this study, we sought to understand users' undirected preferences for interacting with a flexible learning agent in a three-dimensional virtual environment. We chose the digital game Minecraft as it provides an environment that is complex and dynamic yet suitably limited for interactions with a flexible

artificial intelligence (AI) [9]. Based on current research directions and interviews with game developers, we designed a Wizard of Oz prototype of a plausible helpful agent that could operate in Minecraft and learn from user input. We conducted a user study with Minecraft players to observe how they sought to interact with the agent, and to discuss their preferences and concerns for interaction with this type of agent.

This study provides designers of embodied learning agents in virtual worlds with a guide to the affordances that are sought by their users, in the absence of clear signifiers and pre-existing conventions, to assist in matching system design to user expectations [16].

Literature Review

Research interest in human interaction with intelligent agents has increased steadily in the past decade. Of particular relevance is the field of Interactive Machine Learning, which has emerged to study scenarios in which humans act as teachers to learning agents, including both software agents and robots. Studies of human behaviour in these scenarios have consistently shown that users have strong preferences for how to teach agents, which do not always align with the teaching model on which the agent is designed [1]. In particular, users typically focus on giving an agent guidance on how it should behave, in the form of demonstrations and positive prompts, and give relatively little feedback on an agent's past actions, especially negative feedback [1,10,11]. Users are also prone to frustration when required to give repetitive and simplistic input to an agent, which can lead to poor learning outcomes [1]. Fischer et al. found that human teachers would adapt their teaching behaviour to better suit a learning robot based on feedback, but only when

the robot reflected the human's social behaviour (specifically gaze), demonstrating that users apply mental models of robot learning derived from their knowledge of human learning and attention [4]. Similarly, Koenig et al. found that human instructors tended to respond ineffectively to feedback from a robot learner due to a "tendency to map a human-like model onto the capabilities of the robot" [12].

In many cases, the ideal intelligent agent design is either not yet technologically feasible, or prohibitively expensive to produce for a research study. To compensate, HCI researchers have often used the Wizard of Oz method, in which a prototype is secretly operated by a human researcher unbeknownst to the research participant, to test conceptual designs for an intelligent agent or robot [7,15,17]. This has enabled the study of user behaviour with agents to step ahead of the availability of real-world systems. Xu et al. used a Wizard of Oz design to show that users could recognise when an autonomous agent's actions changed, and adapt their own interaction behaviour to suit [20]. Bernotat et al. found that users who were given no specific instructions for how to control a smart home system most often defaulted to speech input [2]. These studies demonstrate that the Wizard of Oz approach is well suited to exploring users' spontaneous or intuitive responses to intelligent agents.

Approach

We recruited 18 participants (aged 11-15, of which 11 were female) from two high schools in the United Kingdom to participate in an observational user study. All participants were required to have played Minecraft in the past. The study consisted of 18 sessions across two weeks, with a single participant in each session.

In each session, the facilitator first interviewed the participant about how often they had played Minecraft, whether they played by themselves or with other people, and which activities and game modes they typically played. The facilitator then asked the participant to complete three simple building tasks in Minecraft. In the first task, the player was given five minutes to build a model boat; this allowed the researcher to observe the player's behaviour in solo Minecraft play. In the second and third tasks, the player was asked to build a maze, and an embodied AI assistant named "help_bot" was introduced.

Virtually embodied intelligent agent design

Help_bot was explained to the player as a prototype AI bot that learned how to act in Minecraft by observing human players' actions. In reality, help_bot was a Wizard of Oz prototype, operated according to a behavioural script by a researcher in another room. Help_bot was given a set of abilities extrapolated from current research directions in machine learning: it "saw" the same visual input as a player; recognised objects within the Minecraft world; and had a limited ability to predict the player's intention, e.g. estimating a larger geometric shape from the placement of a few initial blocks. It was also able to learn from the player's actions, in the form of positive and negative feedback, direct instruction, and labelling (e.g. learning the shape "house" from the text "This is a house").

The base mode of help_bot was to follow and observe the player's avatar. Periodically, the researcher operating help_bot would mentally categorise the player's current action (building, mining/destroying, attacking or waiting/unspecified) and match it. Where possible, help_bot attempted to continue the player's

current task, such as building onto a wall or digging out a pit, and to mine or build with the same block type.

Help_bot also watched for intentional prompts from the player. These included being given a particular material or tool; being hit; or having its recent work reversed, such as when the player destroyed blocks help_bot had recently placed. Each of these prompts was used to update help_bot's model of what it was required to do.

User input conditions

The second and third task in the study involved two conditions for player input to help_bot. The order of these conditions was varied between participants (non-randomly, ensuring a balanced allocation of age, gender and Minecraft experience in each order). In Condition 1, help_bot would respond only to the in-game actions described above. In Condition 2, help_bot would also respond to natural language input, typed through Minecraft's built-in chat channel. Help_bot recognised any instruction that included a reasonably clear action-indicator (verbs such as "build" and "follow") and a stated or implied subject (nouns such as "food" and "me"), and that corresponded with a specified action category. Ambiguous or incomplete instructions prompted a request for clarification from help_bot, in the form of a question mark: "?".

Instructions to participants

Participants were given non-specific instructions on how to interact with help_bot. Participants were told before Condition 1: "Help_bot learns from what it sees you do. You can try to teach it or show it what to do.", and before Condition 2: "Help_bot learns from what it sees you do and what you write in the chat channel. You can try to teach it, show it or tell it what to do. For

example, you could try telling it to bring you something, or ask it to build something." The aim was to observe how participants spontaneously chose to interact with (or ignore) help_bot, as a guide for what kinds of interactions may feel natural and intuitive to users. The instructions did not specify that participants were required to interact with help_bot; one participant chose to ignore help_bot throughout both tasks, and several interacted with it only minimally.

After each task, the facilitator conducted a semi-structured interview with the participant, with questions relating to their thinking during the task, their strategies for guiding help_bot's behaviour, their preferences for interacting with help_bot and how playing with help_bot compared to playing with another human. At the conclusion of the study, participants were informed of the Wizard of Oz nature of help_bot. No participant indicated prior to this debriefing that they suspected help_bot was controlled by a human.

Results

In our analysis of the player session recordings, we identified nine common patterns in the ways players sought to interact with help_bot. We report first on the interactions that used the game controls, and then on interactions that used text input via the chat channel.

Interactions using game controls

The standard game controls consist of mouse and keyboard inputs for movement, selecting and using items, attacking or mining with the selected item, and dropping the selected item. These represent the player's affordances for navigating and interacting with the Minecraft world. Although they were not given specific instructions on how to interact with help_bot,

participants used a consistent set of approaches when attempting to interact with the agent using the game controls. The four common approaches were: *demonstrating, prompting, correcting* and *pointing*.

Demonstrating was the most frequent type of interaction with help_bot. In this interaction, players modelled behaviour they wished help_bot to undertake, such as mining a particular block type or building the initial foundations of a wall. Players were mostly satisfied with the effectiveness of this approach, although they encountered some difficulty in signalling which behaviour was intended to be a demonstration to help_bot and which was not.

Prompting was a less common variation on demonstrating, in which the player used an item or action to suggest a related behaviour. E.g. several players threw help_bot a tool (such as an axe) to indicate that it should mine the type of block that the tool was suited for (wooden blocks). Similarly, a player showed help_bot that they were holding an apple to prompt help_bot to eat its own apple.

Correcting was the second most common type of interaction using the game controls, and usually took the form of reversing the effect of help_bot's recent actions to signal that it should alter its behaviour. Correcting was commonly used as *implicit negative feedback* to refine a behaviour previously initiated through demonstration. E.g. one player began to dig a pit, causing help_bot to follow suit; once help_bot had dug below the desired depth, the player filled in the most recently dug blocks to signal that help_bot should cease digging. *Explicit negative feedback* was rarely

attempted, although in two cases a player hit help_bot to tell the agent it was doing the wrong thing.

Pointing was the final common type of interaction using the game controls. Several players wondered aloud how to direct help_bot's attention to a specific location, e.g. as a designated drop-off point for collected material. In all cases, the eventual solution was to move the avatar to stand on or look at the specific location. In Condition 2, this was typically paired with a text input such as "here". Players expressed that they would prefer a solution that allowed them to precisely indicate locations more quickly and at a distance.

Interactions using text

Participants showed greater variation in their natural language text input compared to their game control input when interacting with help_bot. Grammar varied considerably, with some players writing terse phrases such as "get stone" and others writing full sentences complete with polite speech such as "please" and addressing help_bot by name. Within these variations, however, we identified five main approaches to text-based interaction with help_bot: *instructing, labelling, questioning, encouraging* and *cancelling*.

By far the most frequent type of text input was *instructing*: direct commands for help_bot to perform an action. This was sometimes phrased as a question, as in "can you bring me some wood please", although in almost all such cases a question mark was not used (in contrast to genuine questions). Instructions ranged from simple one- or two-action sets such as "follow me" and "bring me the coal" to more complex concepts such as "build a house". The latter were outside our specifications of what help_bot understood at the start

of the test, but within what it could be taught. Most participants tested one or two such complex requests and fell back on simple requests when these failed.

A few participants used *labelling*: indicating through text that an object or a sequence of actions match a specified term. For example, one participant asked help_bot to “watch”, built a simple hut shape, typed “this is a shelter”, and finally commanded help_bot to “build a shelter”. Several participants suggested labelling as a useful method of automating repetitive work through help_bot, but there was some uncertainty about whether help_bot could accurately judge the boundaries of the object or sequence that was labelled.

Several participants attempted to learn about help_bot through direct *questioning*. Questions were usually, although not always, distinguished from requests by the use of a question mark. Questions referred to help_bot’s inventory (“do you have any wood?”), status (“are you lost?”) and capabilities (“can you fight?”).

Some participants sent *encouraging* messages such as “thank you” and “well done” when help_bot had completed a task. In the post-task interviews, this was explained as a form of positive feedback, to reinforce the behaviour and help the agent learn.

Cancelling messages, such as “stop” and “don’t mine that”, were used when help_bot made categorical errors. E.g. when a participant asked help_bot to “get food”, it attacked a nearby cow rather than looking for fruit, whereupon the player responded with “don’t kill everything”. Cancelling was not used for smaller-scale mistakes, such as placing blocks in the wrong location.

Discussion

The results of our study show broad commonalities in the ways that players approach a virtually embodied intelligent assistant. Although we did not give specific instructions on how to interact with help_bot, our participants followed a small set of interaction patterns, both in text and using the game controls.

Consistent with past studies of non-expert human instructors [1,10], we found that players preferred to teach through demonstration and example rather than explicit feedback. This presents a notable difficulty for interactive machine learning techniques, such as active learning, that rely on explicit user feedback to inform the agent’s learning. Implicit user feedback, in the form of thanks (positive reinforcement) or reversals of the agent’s actions (negative reinforcement), could be used to compensate for a lack of explicit feedback.

Conclusion

This paper contributes an overview of how users spontaneously seek to interact with and teach an embodied learning agent in a virtual world context. Our study lays the groundwork for future study of usability factors for virtually embodied learning agents, and identifies key focus areas for interaction design with learning agents in virtual worlds.

Players’ focus on demonstrating to and observing help_bot reflect patterns of learning behaviour observed in human social groups [6,13,19]. In a future study, we will compare how players behave when they believe they are interacting with an AI and a human, to understand whether aspects of agent design may be informed by theories of human social learning.

Acknowledgements

This work was supported by the Microsoft Research Cambridge Internship Program and the Australian Government Research Training Program Scholarship.

References

1. Saleema Amershi, Maya Cakmak, W Bradley Knox, and Todd Kulesza. 2014. Power to the People: The Role of Humans in Interactive Machine Learning. *AI Magazine* 35, 4.
2. Jasmin Bernotat, Birte Schiffhauer, Friederike Eyssel, et al. 2016. Welcome to the Future – How Naïve Users Intuitively Address an Intelligent Robotics Apartment. In A. Agah, J.-J. Cabibihan, A.M. Howard, M.A. Salichs, and H. He, eds., *Social Robotics: 8th International Conference, ICSR 2016, Kansas City, MO, USA, November 1-3, 2016 Proceedings*. Springer International Publishing, Cham, 982–992.
3. Roland Buchner, Daniela Wurhofer, Astrid Weiss, and Manfred Tscheligi. 2013. Robots in Time: How User Experience in Human-Robot Interaction Changes over Time. *Proceedings of the 5th International Conference on Social Robotics*, Springer, 138–147.
4. Kerstin Fischer, Katrin S Lohan, Chrystopher Nehaniv, and Hagen Lehmann. 2013. Effects of Different Kinds of Robot Feedback. *Proceedings of the 5th International Conference on Social Robotics*, Springer, 260–269.
5. Susanne Frennert, Håkan Efring, and Britt Östlund. 2013. What Older People Expect of Robots: A Mixed Methods Approach. *Proceedings of the 5th International Conference on Social Robotics*, Springer, 19–29.
6. James Paul Gee and Judith L. Green. 1998. Discourse Analysis, Learning, and Social Practice: A Methodological Study. *Review of Research in Education* 23: 119–169.
7. Michael A. Goodrich and Alan C. Schultz. 2008. Human–Robot Interaction: A Survey. *Foundations and Trends® in Human–Computer Interaction* 1, 3: 203–275.
8. Kristina Höök. 2000. Steps to Take Before Intelligent User Interfaces Become Real. *Interacting with Computers* 12, 4: 409–426.
9. Matthew Johnson, Katja Hofmann, Tim Hutton, and David Bignell. 2016. The Malmo Platform for Artificial Intelligence Experimentation. *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-16)*, 4246–4247.
10. Tasneem Kaochar, Raquel Torres Peralta, Clayton T. Morrison, Ian R. Fasel, Thomas J. Walsh, and Paul R. Cohen. 2011. Towards Understanding How Humans Teach Robots. *Proceedings of the 19th International Conference on User Modeling, Adaption, and Personalization*, Springer-Verlag, 347–352.
11. W. Bradley Knox, Brian D. Glass, Bradley C. Love, W. Todd Maddox, and Peter Stone. 2012. How Humans Teach Agents. *International Journal of Social Robotics* 4, 4: 409–421.

12. Nathan Koenig, Leila Takayama, and Maja Matarić. 2010. Communication and Knowledge Sharing in Human–Robot Interaction and Learning from Demonstration. *Neural Networks* 23, 8–9: 1104–1112.
13. Jean Lave and Etienne Wenger. 1991. *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press.
14. Pattie Maes. 1994. Agents That Reduce Work and Information Overload. *Communications of the ACM* 37, 30–40.
15. David Mulsby, Saul Greenberg, and Richard Mander. 1993. Prototyping an Intelligent Agent through Wizard of Oz. *Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems*, ACM, 277–284.
16. Donald A. Norman. 1994. How Might People Interact with Agents. *Communications of the ACM* 37, 7: 68–71.
17. Laurel Riek. 2012. Wizard of Oz Studies in HRI: A Systematic Review and New Reporting Guidelines. *Journal of Human-Robot Interaction* 1, 1: 119–136.
18. Ben Shneiderman. 1982. The Future of Interactive Systems and the Emergence of Direct Manipulation. *Behaviour & Information Technology* 1, 3: 237–256.
19. Gerry Stahl. 2013. Theories of Cognition in Collaborative Learning. In C.E. Hmelo-Silver, ed., *The International Handbook of Collaborative Learning*. Routledge, New York, NY, USA, 74–90.
20. Yong Xu, Yoshimasa Ohmoto, Kazuhiro Ueda, et al. 2011. Active Adaptation in Human-Agent Collaborative Interaction. *Journal of Intelligent Information Systems* 37, 1: 23–38.