Player expectations of a learning AI companion in Minecraft

Fraser Allison
Microsoft Centre for Social NUI
The University of Melbourne
fraser.allison@unimelb.edu.au

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

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

Keywords
Machine learning, Minecraft, AI, human-agent interaction, Wizard of Oz

ABSTRACT

In recent years, machine learning techniques have achieved remarkable success in numerous fields of application, not least in digital games. Games have been widely used as testing grounds and performance benchmarks for machine learning-based artificial intelligence (AI) (Yannakakis & Togelius, 2015), from traditional board games (Silver et al., 2016) to real-time arcade games (Mnih et al., 2015; Shaker et al., 2013) and complex three-dimensional virtual worlds that more closely approximate physical space, such as Minecraft (Johnson, Hofmann, Hutton, & Bignell, 2016). However, commercial game developers have been slow to take up machine learning techniques for in-game character AI, due in part to concerns that learning-based AI would be difficult for players to understand, or simply not fun to play with (Muñoz-Avila, Bauckhage, Bida, Congdon, & Kendall, 2013; Yannakakis & Togelius, 2015). This study seeks to address these concerns by investigating players' expectations and preferences for interacting with a friendly in-game character that learns from player input.

We conducted an observational study on 18 participants (aged 11-15), who undertook building tasks in Minecraft with the assistance of a robotic-looking character named "help_bot". Help_bot was introduced as a prototype AI that learned from observing the player's actions and instructions. In reality, help_bot was a Wizard of Oz prototype, controlled from another room by a hidden researcher according to a behavioural script. In one task condition, help_bot responded only to player actions. In another task condition, help_bot also responded to written natural language prompts. Before and after each task, the research facilitator conducted semi-structured interviews with the player, which asked about their past Minecraft experience; their thinking during the task; their strategies for engaging with help_bot; and their opinions and preferences towards help_bot. Players were informed of the Wizard of Oz nature of help_bot at the end of the study period.

The primary research objective was to observe how players spontaneously chose to interact with this novel learning-based character under minimal instruction, to determine

Extended Abstract Presented at DiGRA 2017

what modes of interaction were most intuitive to or expected by players. In the post-task interviews, players reported that they drew their expectations about help_bot from several sources: their knowledge of friendly villagers and pets in Minecraft; online multiplayer social conventions; real world social conventions; Minecraft console commands ("cheats"); programming languages; Hollywood film depictions of AI; and human cognition. These informed a range of player strategies for engaging with help_bot, some of which were well suited to help_bot's nominal abilities and others that were not.

In a thematic analysis of player actions and interview responses, we noted two main dimensions along which the style of interaction with help_bot varied between players. The first and most obvious was the level of engagement. While most players divided their attention between help_bot and the building task they had been assigned, several players demonstrated minimal interest in help_bot, and one ignored help_bot entirely. Conversely, several players were highly engaged with help_bot throughout the study, abandoning the building tasks in favour of testing the limits of what they could get help_bot to do.

The second dimension of player variability was the degree to which players anthropomorphised help_bot as a conscious character. Those who talked about help_bot as though it was a conscious character were typically more polite in their text inputs to help_bot, expressed more empathy when help_bot was attacked or became stuck, and expected a higher degree of reasoning ability. In the interviews, they were more likely to say that they would enjoy having help_bot as a companion character. Players at the other end of the scale, who took a more instrumental view of help_bot, expressed less empathy for the character, were terser in their text inputs, and expected less reasoning ability. They were more likely to say they would use help_bot to automate boring tasks in the game.

Notably, the engagement and anthropomorphisation dimensions were at least partially independent of each other. Some players who were unengaged with help_bot nevertheless treated it politely as a character, and conversely some players who anthropomorphised help_bot little were highly engaged with maximising its usefulness as a tool for batch-processing repetitive actions.

These differences in players' thinking about help_bot illustrate a challenge for developers of intelligent game companions. Different expectations and mental models of how an AI learns and "thinks" lend themselves to different applications of learning AI in games. Being able to cue an appropriate mental model of a learning AI will help to ensure players accurately understand the constraints and affordances of the game system, and preserve a game experience that accords with the designers' intentions.

PRESENTING AUTHOR BIO

Fraser Allison is a PhD candidate in human-computer interaction at the University of Melbourne, based in the Microsoft Research Centre for Social Natural User Interfaces. His research focuses on voice interfaces, intelligent agents and digital games.

ACKNOWLEDGMENTS

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

BIBLIOGRAPHY

Johnson, M., Hofmann, K., Hutton, T., & Bignell, D. (2016). The Malmo Platform for Artificial Intelligence Experimentation. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-16)* (pp. 4246–4247).

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., … Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, *518*(7540), 529–533. https://doi.org/10.1038/nature14236

Muñoz-Avila, H., Bauckhage, C., Bida, M., Congdon, C. B., & Kendall, G. (2013). Learning and Game AI. In S. M. Lucas, M. Mateas, M. Preuss, P. Spronck, & J. Togelius (Eds.), *Artificial and Computational Intelligence in Games* (Vol. 6, pp. 33–43). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

Shaker, N., Togelius, J., Yannakakis, G. N., Poovanna, L., Ethiraj, V. S., Johansson, S. J., … Gallagher, M. (2013). The Turing Test Track of the 2012 Mario AI Championship: Entries and evaluation (pp. 1–8). Presented at the IEEE Conference on Computational Intelligence in Games (CIG), IEEE. https://doi.org/10.1109/CIG.2013.6633634

Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., … Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, *529*(7587), 484–489. https://doi.org/10.1038/nature16961

Yannakakis, G. N., & Togelius, J. (2015). A panorama of artificial and computational intelligence in games. *IEEE Transactions on Computational Intelligence and AI in Games*, *7*(4), 317–335.

Proceedings of DiGRA 2011 Conference: Think Design Play.