

Spotting the Difference: Identifying Player Opponent Preferences in FPS Games

David Conroy, Peta Wyeth, and Daniel Johnson

Queensland University of Technology,
Science and Engineering Faculty,
Brisbane, Australia
{david.conroy,peta.wyeth,dm.johnson}@qut.edu.au

Abstract. This paper describes a study designed to understand player responses to artificially intelligent opponents in multi-player First Person Shooter games. It examines the player's ability to tell the difference between artificially intelligent opponents and other human players, and investigates the players' perceptions of these opponents. The study examines player preferences in this regard and identifies the significance of the cues and signs players use to categorise an opponent as artificial or human.

Keywords: Video Games, Artificial Intelligence, Identification Processes, User Study, Multi-player.

1 Introduction

Artificially intelligent (AI) opponents or bots¹ are an important feature of First Person Shooter (FPS) video games. There have been many advances in the field of bot AI technology, particularly in relation to attempting to better mimic player-like behaviour. Despite these advances, there are still unanswered questions about what players expect from FPS AI and how bots impact on the player experience. The FPS genre of videogames can be defined as a virtual world in which the player moves in a first-person perspective to achieve the goals of the game. These game goals can vary depending on the different kinds of play or combat mode [1][2]. In the game mode dubbed 'deathmatch', the objective is to gain points by killing or 'fragging' other opponents in a free-for-all style. During a match, a player will interact with other opponents who can be controlled by other human players or by the game itself through AI controlled bots. The types of games that are the focus of this research are these deathmatch style games.

The study detailed in this paper forms part of a program of research designed to investigate claims that current AI designs are unrealistic, predictable and exploitable and influence game play negatively [3]. The current paper examines player preferences for in-game opponents. It explores the contrast between the behaviour of real

¹ Bots – computer controlled agents that mimic player behaviours performing game play tasks.

players and AI controlled bots in the game. The focus of the study is to better understand players' views through:

- Investigating the extent to which players are able to identify human players and bots in multiplayer environments; and
- Discerning which opponent behaviours are the most useful in assisting players identify bots from human opponents.

This paper is not an analysis of the algorithms underlying the control of bot AI systems in FPS games or an examination of the psychological characteristics of participants that make them perceive AI agents the way they do. The research is designed to understand how readily a player can identify bots from humans and the types of opponent behavior traits they use in this identification.

2 Background

There is evidence that players have different experiences playing against computer controlled opponents as opposed to competing against human players, especially with respect to presence, flow and enjoyment [4]. Results show that most people generated higher physiological arousal and more fun while playing against another human. It may be that the exploitive and repetitive nature of NPCs in games reduces player enjoyment and causes them to look for more challenging game play experiences [5]. Interviews have identified that, while players appreciate challenging bots that give players a chance to practice and improve their skills, they also flag negative aspects such as unrealistic behavior, unexpected behaviour (e.g. not taking cover or retreating), problems with coordination and unfair game advantages [6]. "Subhuman" behaviour, behaving in predictable ways that made it easy for human players to ascertain the weaknesses of bots and invent unconventional ways to defeat them, was identified as a particular weakness of opponent AI.

Human control of a game avatar and bot controlled AI are inherently different. Players do not naturally have a fully intuitive sense of orientation and action in virtual environments and must invest time and energy to master the control interface and learn the mechanics of games [7]. They require time to master game actions, e.g. shooting at targets [8]. Conversely, bot AI often includes a notion of shooting inaccurately through the random 'missing' of targets. This process doesn't consider improvement or progression in expertise. When expert judges were asked to identify the human player from the game play of a human and a bot in a one-on-one deathmatch environment, little correlation between skill levels and ratings of humanness was found [1]. Such results might call into question the necessity to create bot AI with multiple skill levels (e.g. [9]). Dynamic game systems that change the priority of a tactic depending on success ratings might be more akin to the play style of human players, especially in fast paced environments such as FPS games [5]. Research has demonstrated that it might be useful to consider bot play-styles based on survival as this play style has been shown to improve performance of bots [2] and strongly rewards human players as well [9].

Previous studies that have examined players' categorisation of human and bot opponents have demonstrated that it is not necessarily easy to judge accurately. Hings-ton [1] determined that people found it difficult when asked to observe a human and a bot in a one-on-one deathmatch environment and rate the likeness to the play style to a human player. Judges were often far from accurate. High aggression, adaptability and tactical game play were the characteristics that judges most easily recognised as human. Missing behaviours, stupid behaviours and low aggression also indicated bot-like game play. A study involving functional magnetic resonance imaging (fMRI) found higher levels of brain activity when participants believed they were competing against other human participants as opposed to a bot, even if in actual fact they were not [10]. Participant's also commented they were convinced of the alleged opponent's humanity and found those games to be more compelling.

From the literature it is clear that there are still unanswered questions related to players' abilities to identify bots from humans, and their perceptions of both. It is equally clear that we need this understanding to guide the future development of bot AI, to ensure that a positive player experience results.

3 Method

A user study was conducted to examine players' responses to bot and human opponents. The study consisted of two components; a multi-player gameplay session and a participant questionnaire. Five groups of four participants were involved, resulting in total of 20 participants taking part in the user study. The game used for the study was Quake III: Arena by iD games. The Quake III game client was used mainly for its ease of modification and setting up of dedicated multiplayer games.

3.1 Multiplayer Game Play Sessions

Each multiplayer game play session required four participants to play four multiplayer death matches consecutively with a brief questionnaire acting as intermission between games. Participants were positioned in a room with dividers blocking their view of the other participants. Participants were briefed on set-up procedures and were provided with simple step by step instructions. Each participant was required to wear headphones to eliminate vocal noises emitted by other participants. For each participant, two of the four multiplayer games were against a single opponent (1v1). Each of these games lasted approximately seven minutes. The other two games consisted of three opponents, making them four player matches (1v1v1v1). Each of these games lasted approximately 10 minutes. Two of the four games, one 1v1 and one 1v1v1v1, had bot opponents for the players to compete against. The order of presentation of the first two 1v1 games was randomized for opponent types. The 1v1v1v1 game that contained bot opponents had two bots and two human player. By the end of the study session each participant would have competed against three bots and every other participant in the room at least once.

3.2 Participant Questionnaire

The participant questionnaires consisted of a series of questions that players answered prior to playing any games, immediately after playing specific game session types and finally, after all games were completed. Questions asked before playing related to player demographic information. After each multiplayer session, participants rated the humanness level of all opponents on a Likert scale between 1 and 7, with 1 definitely being a bot, 4 being neutral/undecided and 7 definitely being a human. Any participant who rated a bot between 1 and 3 (being a bot) could be described as correctly rating the opponent as a bot. If they rated the bot from 5 to 7, they have incorrectly rated the bot as a human. Anyone who felt they could not decide rated the opponent a 4 and were therefore neutral about their decision. This same process is applied to when the opponent was a human, but in reverse.

After the first two 1v1 matches, players were queried on which opponent they preferred (Opponent 1 or Opponent 2). After all games were completed, participants completed a final section of the questionnaire regarding how useful they felt certain behaviours were in identifying whether an opponent was a human or bot. This consisted of a list of 11 behaviours (e.g. dodging, aiming) and another Likert scale to rate the extent to which each behaviour assisted in identifying a bot from a human. A 5 point scale was incorporated for this assessment, with 1 indicating that a behaviour failed to assist in classification and 5 indicating that it provided strong evidence to guide a player's decision. The 11 behaviours chosen for this identification process were actions found through previous research [3][11] to be most prevalent in participants playing FPS games. These behaviors were extracted and analysed from previous research data using an adapted time and event sampling technique [12]. Participants were not permitted to communicate with the other participants in any way either during game play or when completing the questionnaire.

4 Results

4.1 Two Player Participant Data

The 1v1 data allowed for a direct comparison of each participant's ratings after playing a single bot and a single human player. Each of the 20 participants only competed against one bot and one human in 1v1 games, making 20 ratings for each opponent type (bot and human). The results can be found in Fig. 1.




Fig. 1. 1v1 Participant Assessment Totals: green – correct, orange – neutral and red – incorrect

The results from the 1v1 participant data in Fig. 1 show that players are able to correctly categorise a bot and a human in a majority of instances, with only a small number of participants maintaining a neutral stance. Participants more accurately identified humans correctly. Participants were generally more definite in their assessment when categorising an opponent as human – 55% as “definite” or “most likely” as opposed to 40% for bots. The number of incorrect identifications made in the 1v1 games appears to be similar between both the bot and human games. However, while 3 participants recorded the bot as "definitely a human", no participants rated their human opponent as "definitely a bot".

4.2 Four Player Participant Data

The four player deathmatch games provided data that allowed for comparison between participant ratings of human and bot opponents in games with multiple opponents. Each participant played one match against 1 human and 2 bot opponents and one match against 3 human opponents. This resulted in a total of 40 possible bot identifications and 80 total human identifications across all the participants. Results are displayed in Fig. 2.




Fig. 2. Four Player Participant Assessment Totals

The data from Fig. 2 reveals that people can correctly identify bots from humans more often than not. However, the percentages for both opponent types decreased in comparison to the 1v1 data, with a larger number of participants choosing incorrectly. Compared to the 1v1 participant data in Fig. 1, there is a significant increase in the number of participants who were indecisive or neutral in their decision. Many participants did not rate all 3 opponents as human in the final 4 player human game sessions. While most participants did manage to rate human opponents on the correct side of the scale (i.e. a rating between 5 and 7), some participants believed there to be at least one other bot in the game, if not two, or could not decide (see Table 1). There were more examples of participants gravitating towards an indecisive vote (i.e. a rating of 3 or 4) for the four player, all-human games. Thirteen of the 18 incorrect assessments of humans as bots, occurred in the 4 player all-human game sessions. The three examples where participants assessed human opponents as “definitely a bot” also occurred in these sessions.

Table 1. Final Four Player Game Participant Ratings

Bot->Human Rating Scale	1	2	3	4	5	6	7
Human Opponents (60)	3 (5%)	6 (10%)	4 (6.6%)	11 (18.3%)	12 (20%)	12 (20%)	12 (20%)

4.3 Opponent Preferences

Information about opponent preferences was gathered from 1v1 gameplay scenarios. A question asked participants which opponent they preferred. Participants were not aware of whether an opponent was a human or a bot. Seventeen participants chose the opponent who was human while only three chose the opponent that was a bot. Interestingly, most participants preferred the opponent who was human, even if they categorised that opponent as a bot. Two participants indicated that they preferred the human opponent, yet perceived it as a bot. The three participants who preferred the bot opponent, perceived the opponent as a human.

4.4 Opponent Behaviour Identification Data

The final section of the questionnaire was used to examine which opponent behaviours are most useful in assisting players in identifying bots from human opponents. The mean and standard deviation of this data is displayed in Fig. 3. The data was analysed to identify which behaviours were deemed most important in identifying human opponents from bot opponents. Behaviours identified as most important in the assessment process, with high mean and mode scores, were detecting, aiming, camping, response to player, fleeing and pursuing.




Fig. 3. Opponent Behaviour Identification Data

5 Discussion

The data concerning player's preference in opponent type yielded interesting results. In the 1v1 games, participants preferred the human opponent to the bot at a ratio of 17:3. There were two participants who perceived both their 1v1 games as having bot opponents, yet indicated their preference for the game where the opponent was a human. From these results it appears that players have a greater preference for human opponents; even those participants who chose bots, perceived them as human. The level of accuracy in classifying bot opponents (60% for 1v1 games and 45% for 4

player games) was lower than expected. For the four player games, this reduction in accuracy may be attributable to more participants being unsure, with neutral assessments increasing in count. While accuracy in determining human players also decreased in four player scenarios (70% in 1v1 versus 60% in 1v1v1v1), the decrease was not as marked. This may indicate that players became more adept at identifying characteristics of human opponents while bot-like activity remains ambiguous for longer.

Players were consistently better at determining human opponents from bot opponents. They also did so with more confidence in the 1v1 scenarios. The data indicates that game opponents are consistently harder to classify if they are bots. Accuracy differences of 10% in 1v1 games and 15% in four player games needs to be examined. When uncertain within the bot scenario, participants appear to be more likely to guess human. When uncertain in the human scenario, they may be more likely to select human. A possible interpretation is that players expect bots to behave in certain “bot-like” ways and when bots don’t meet these expectations, players classify them as human. Consequently, further exploration of “bot-like” behavior is necessary in order to design bots that meet player expectations and have the appeal of human opponents.

Results related to the behaviours deemed most useful for opponent type were informative. The data presented in Fig. 3 shows that behaviours related to the reactions an opponent has to player actions are more useful in determining bots from humans. From the data, detecting, camping, aiming/inaccuracy, pursuing, fleeing and general response to player actions were found to be the major contributors to opponent identification. These behaviours have a higher percentage of participants determining them useful for identifying opponents with little variation in distribution across the behaviour identification rating scheme.

6 Conclusions and Future Work

The research described in this paper was conducted to support a larger research program. The current study examines the players’ ability to tell the difference between bot and human opponents based on exhibited behaviours. This user study was designed to understand the situation from the player’s point of view. While the study revealed some unexpected results, data analysis supports much of our existing research, especially with respect to player preferences for human opponents. Players were capable of identifying bots from humans more often than not, yet with slightly less capability than expected. This shows identification is not straightforward, and the reasoning around decisions, while generally indicating that reaction and response are the key indicators, needs further exploration.

The data generated from the existing study has revealed information which can be applied to future research. In particular, the data describing which behaviours both bots and humans perform that reveal their true identity has provided a stronger argument for the much larger research study. To illustrate with example, the bot’s response to the player was deemed a better predictor than correct weapon usage, or jumping. This is not to say that the way bots jump in FPS games is human-like, it

simply implies that the act of jumping does not act as a clear predictor. Response to the player's presence on the other hand appears to be a good predictor. Investigation of which behaviours have been used by those participants who accurately identified humans from bots is a logical next step. From this data it is possible to develop a picture of how one might better design AI behaviours that players find more engaging. It supports the notion of modifying bot AI to better respond to player actions, particularly concerning attacking, retreating, fleeing and hiding. This information will be integrated into the expanding program of research currently being undertaken. Future work will expand on this study, to develop a model of bot AI that players find more interesting and engaging to compete against.

References

1. Hingston, P.: A Turing Test for Computer Game Bots. *IEEE Transactions on Computational Intelligence and AI in Games* 1, 169–186 (2009)
2. Esparcia-alc, A.I., Member, S., Mora, A., Merelo, J.J., Garc, P.: Controlling bots in a First Person Shooter Game using Genetic Algorithms. *Environment* (2010)
3. Conroy, D., Wyeth, P., Johnson, D.: Modeling player-like behavior for game AI design. In: *Proceedings of the 8th International Conference on Advances in Computer Entertainment Technology - ACE 2011*, p. 1. ACM Press, New York (2011)
4. Weibel, D., Wissmath, B., Habegger, S., Steiner, Y., Groner, R.: Playing online games against computer- vs. human-controlled opponents: Effects on presence, flow, and enjoyment. *Computers in Human Behavior*. 24, 2274–2291 (2008)
5. Hartley, T.P., Mehdi, Q.H.: In-game tactic adaptation for interactive computer games. In: *16th International Conference on Computer Games (CGAMES 2011)*, pp. 41–49 (2011)
6. Clarke, D., Duimering, P.R.: How computer gamers experience the game situation: a behavioral study. *Computers in Entertainment* 4, 6 (2006)
7. Przybylski, A.K., Rigby, C.S., Ryan, R.M.: A motivational model of video game engagement. *Review of General Psychology* 14, 154–166 (2010)
8. Rayner, C.: Player Modelling for Cursor-Driven Games. *Challenges* (2007)
9. Laird, J.E.: Using a computer game to develop advanced AI. *Computer* 34, 70–75 (2001)
10. Krach, S., Blümel, I., Marjoram, D., Lataster, T., Krabbendam, L., Weber, J., van Os, J., Kircher, T.: Are women better mindreaders? Sex differences in neural correlates of mentalizing detected with functional MRI. *BMC Neuroscience* 10, 9 (2009)
11. Conroy, D., Wyeth, P.: Building Better Bad Guys: A New Framework for Game AI Design. In: *Proceedings of the 7th Australasian Conference on Interactive Entertainment, IE 2010*, pp. 0–2. ACM Press (2010)
12. Bushnell, D., Irwin, M.: *Observational Strategies for Child Study*. Rinehart and Winston, Holt (1980)