Using Reinforcement Learning to Create Engaging AI Bot Behaviour in the FPS Genre

George Blackwell

School of Computing & Engineering, University of Huddersfield

Introduction

New and interesting AI behaviour is becoming increasingly desirable within the video game industry. This is especially true when developing competitive AI agents for the first-person shooter genre. Which has been a cornerstone of the industry since its emergence. This project will employ neural networks and genetic algorithms to evolve various intelligent AI agents through reinforcement learning. To evaluate their effectiveness, they will be trained within a simple game of 'Team Deathmatch' against a team of conventional scripted AIs. The trained AIs will operate and train independently and take various inputs from their environment, as a Human player would.

Project Aims

- Identify and segment various areas of the bot's behaviour into separate trainable tasks.
- Develop suitable **feedforward neural networks** for each trainable task.
- Train the bots using a **genetic algorithm** as a **reinforcement learning** (RL) technique.
- Produce interesting behaviour that is able to compete with scripted bots.

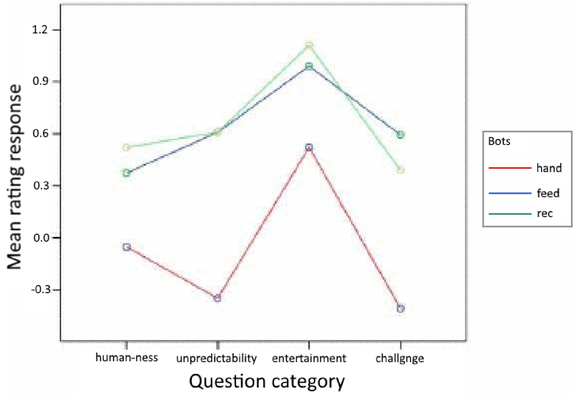

Figure 1: Feedback given by human players on the 'enjoyment' level of various bot types. The trained bots were perceived to act more 'human' and thus more interesting [1].

Traditional Bots

Traditional FPS bots can differ significantly from game to game. However, they generally follow the same traditional rule set of navigating the map, engaging in combat and using/picking up weapons and items. These rules are usually defined through scripted and hard-coded state machines. Even with years of development, traditional AI bots are likely to exhibit undesirable behaviour. This undesirable behaviour can manifest itself as an exploitable quirk, which players can abuse to sidestep difficult encounters. Alternatively, robust but simple behaviour is also an issue. These unwanted or dull behaviours are often masked with unfair adjustments, such as increased health or damage.

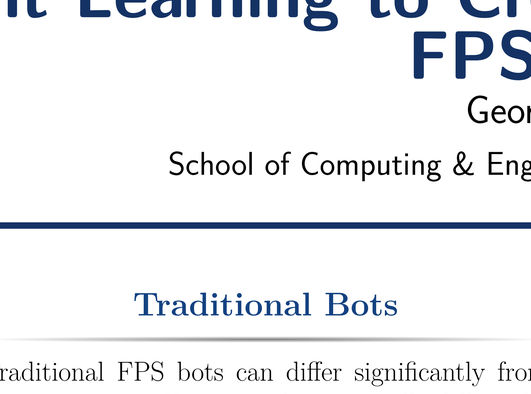

Figure 2: Quake and Unreal have been used extensively within the machine learning research field.

Why Reinforcement Learning?

Compared to other machine learning approaches, RL is the most suited for developing complex bot behaviour within an FPS game. Just as a human player would, the learning bot is interacting with its environment and receiving feedback based on its actions. The feedback can be in the form of a reward for a good action or a punishment for a bad action. The bot's behaviour is then improved based on the rewards and punishments it receives. This method of learning resembles the trial-and-error approach humans take to learning and complements complex game environments [2].

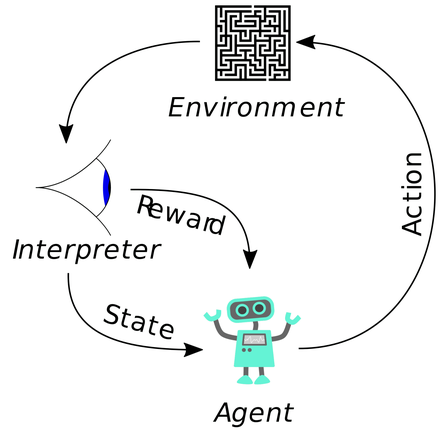

Figure 3: A simple flowchart explaining how reinforcement learning functions.

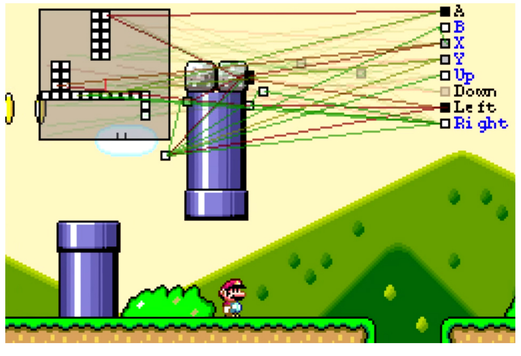

Figure 4: MarI/O, a program that employs a genetic algorithm to teach a computer to play Mario without any assistance.

Are RL Trained Bots Viable?

If executed correctly, RL trained bots could offer the player a refreshing experience that exchanges inflated enemy health and overly precise movement for engaging 'human-like' behaviour. Players are attracted to online player verse player games because there is an element of unpredictability that human players offer and conventional bots don't. However, RL is not without its disadvantages. Most notably, the issue of scaling with increased complexity. A common method to solve this issue is to break the problem down into smaller tasks and then tackle them individually. This then presents the new task of combining them into a suitable behaviour set.

References

[1] Bhuman Soni and Philip Hingston.
Bots trained to play like a human are more fun.
In *Neural Networks, 2008. IJCNN 2008.(IEEE World Congress on Computational Intelligence). IEEE International Joint Conference on*, pages 363–369. IEEE, 2008.

[2] Hao Wang, Yang Gao, and Xingguo Chen.
Rl-dot: A reinforcement learning npc team for playing domination games.
IEEE Transactions on Computational intelligence and AI in Games, 2(1):17–26, 2010.

University of HUDDERSFIELD