

AI in human–computer gaming: Techniques, challenges and opportunities

December 29 2023



The corresponding AIs cover AlphaGo, AlphaGo Zero, AlphaZero, Libratus, DeepStack, DouZero, Suphx, FTW, AlphaStar, OpenAI Five, JueWu1, and Commander. Credit: Beijing Zhongke Journal Publising Co.

Human–computer gaming has a long history and has been a main tool for verifying key artificial intelligence technologies. The Turing test, proposed in 1950, was the first human–computer game to judge whether a machine has human intelligence. This has inspired researchers to develop AI systems (AIs) that can challenge professional human players.

A typical example is a checkers AI called Chinook, which was developed in 1989 to defeat the <u>world champion</u>. The target was



achieved when it beat Marion Tinsley in 1994. Later, Deep Blue from IBM beat chess grandmaster Garry Kasparov in 1997, setting a new era in the history of human–computer gaming.

In recent years, researchers have witnessed the rapid development of human–computer gaming AIs, from the DQN agent, AlphaGo, Libratus, and OpenAI Five to AlphaStar. These AIs can defeat professional human players in certain games with a combination of modern techniques, indicating a big step in the decision-making intelligence.

For example, AlphaGo Zero, which uses Monte Carlo tree search, selfplay, and <u>deep learning</u>, defeats dozens of professional go players, representing powerful techniques for large state perfect information games. OpenAI Five, using self-play, deep reinforcement learning, and continual transfer via surgery, became the first AI to beat the world champions at an eSports game, displaying useful techniques for complex imperfect information games.

After the success of AlphaStar and OpenAI Five, which reach the professional human player level in the games StarCraft and Dota2, respectively, it seems that current techniques can solve very complex games. Especially the breakthrough of the most recent human–computer gaming AIs for games such as the Honor of Kings and Mahjong obeys similar frameworks of AlphaStar and OpenAI Five, indicating a certain degree of universality of current techniques.

So, one natural question arises: What are the possible challenges of current techniques in human–computer gaming, and what are the future trends? A new paper <u>published</u> in *Machine Intelligence Research* aims to review recent successful human–computer gaming AIs and tries to answer the question through a thorough analysis of current techniques.

Based on the current breakthrough of human-computer gaming AIs



(most published in journals such as *Science* and *Nature*), researchers survey four typical types of games, i.e., board games with Go; card games such as heads-up no-limit Texas hold'em (HUNL), DouDiZhu, and Mahjong; first person shooting games (FPS) with Quake III Arena in capture the flag (CTF); real-time strategy games (RTS) with StarCraft, Dota2, and Honor of Kings. The corresponding AIs cover AlphaGo, AlphaGo Zero, AlphaZero, Libratus, DeepStack, DouZero, Suphx, FTW, AlphaStar, OpenAI Five, JueWu, and Commander.

The remainder of the paper is organized as follows. In Section 2, researchers describe the games and AIs covered in this paper. Based on the recent progress of human–computer gaming AIs, this paper reviews four types of games and their corresponding AIs, i.e., board games, card games, FPS games, and RTS games. To measure how hard a game is to develop professional human-level AI, researchers extract several key factors that challenge intelligent decision-making, which are imperfect information, long time horizon, in-transitive game and multi-agent cooperation.

Section 3 is about board game AIs. The AlphaGo series is built based on Monte Carlo tree search (MCTS), which is widely utilized in previous Go programs. AlphaGo came out in 2015 and beats European Go champion Fan Hui, which was the first time that an AI won against professional players in a full-size game, Go without Renzi. Afterward, an advanced version called AlphaGo Zero was developed using different learning frameworks, which needs no prior professional human confrontation data and reaches superhuman performance. AlphaZero uses a similar learning framework to AlphaGo Zero and explores a general reinforcement learning algorithm, which masters Go along with another two <u>board games</u>, chess, and Shogi.

Section 4 introduces card game AIs. The Card game, as a typical inperfect information game, has been a long-standing challenge for



artificial intelligence. DeepStack and Libratus are two typical AI systems that defeat professional poker players in HUNL. They share the same basic technique, i.e., counterfactual regret minimization (CFR). Afterward, researchers are focusing on Mahjong and DouDiZhu, which raise new challenges for artificial intelligence. Suphx, developed by Microsoft Research Asia, is the first AI system that outperforms most top human players in Mahjong. DouZero, designed for DouDiZhu, is an AI system that was ranked first on the Botzone leaderboard among 344 AI agents.

First-person shooting game AIs are shown in Section 5. CTF is a typical three-dimensional multiplayer first-person video game in which two opposing teams are fighting against each other on indoor or outdoor maps. The settings for CTF are very different from the current multiplayer video games. More specifically, agents in CTF cannot access the state of other players, and agents in a team cannot communicate with each other, making this environment a very good testbed for learning agents to emerge communication and adapt to zero-shot generation. Zero-shot means that an agent cooperated or confronted is not the agent trained, which can be human players and arbitrary AI agents. Based solely on pixels and game points like a human as input, the learned agent FTW reaches a strong human-level performance.

Section 6 is about RTS game. The RTS game, as a typical kind of video game, with tens of thousands of people fighting against each other, naturally becomes a testbed for human–computer gaming. Furthermore, RTS games usually have a complex environment, which captures more of the nature of the real world than previous games, making the breakthrough of such games more applicable. AlphaStar, developed by DeepMind, uses general learning algorithms and reaches the grandmaster level for all three races for StarCraft, which also outperforms 99.8% of human players who are active on the European server. Commander, as a lightweight computation version, follows the same learning architecture



as AlphaStar, which uses an order of magnitude less computation and beats two grandmaster players in a live event. OpenAI Five aims to solve the Dota2 game, which is the first AI system to defeat the world champions in an eSports game. As a relatively similar eSports game to Dota2, Honor of Kings shares most similar challenges, and JueWu becomes the first AI system that can play full RTS games instead of restricting the hero pool.

In Section 7, researchers summarize and compare the different techniques utilized. Based on the current breakthrough of human-computer gaming AIs, currently utilized techniques can be roughly divided into two categories, i.e., tree search (TS) with self-play (SP) and distributed deep reinforcement learning (DDRL) with self-play or population-play (PP). It should be noted that researchers just mention the basic or key techniques in each category, based on which different AIs usually bring in other key modules based on the games, and those new modules are sometimes not generic across games. The tree search has two kinds of representative algorithms: MCTS, usually used for perfect information games, and CFR, conventionally designed for imperfect information games. As for population-play, it is used for three situations: different players/agents do not share the same policy network due to the game characteristics; populations can be maintained to overcome the game theoretical challenges such as non-transitivity; populations combined with population-based training to learn scalable agents. With the comparison, researchers discuss two points as follows: How to reach Nash equilibrium, and how to become general technology.

In Section 8, researchers show the challenges in current game AIs, which may be the future research direction of this field. Even though big progress has been made in human–computer gaming, current techniques have at least one of three limitations. Firstly, most AIs are designed for a specific human–computer game or a map of a specific game, and the AIs learned are not able to be used even for different maps of a game.



Moreover, not enough experiments are performed to validate the AI's ability when a disturbance is brought into the <u>game</u>. Secondly, training the above AIs requires a large number of computation resources. Due to the huge hardware resource threshold, only a limited number of organizations are capable of training high-level AIs, which will obstruct most scientific research from an in-depth study of the problem. Thirdly, most AIs are evaluated based on their winning ability against limited professional human players, and a claim of reaching the expert level may be a little exaggerated. The potential directions and challenges faced by the above limitations are presented in this part.

This paper summarizes and compares techniques of current breakthroughs of AIs in human–computer gaming. Through this survey, researchers hope that beginners can quickly become familiar with the techniques, challenges, and opportunities in this exciting field, and researchers on the way can be inspired for deeper study.

More information: Qi-Yue Yin et al, AI in Human-computer Gaming: Techniques, Challenges and Opportunities, *Machine Intelligence Research* (2023). DOI: 10.1007/s11633-022-1384-6

Provided by Beijing Zhongke Journal Publising Co. Ltd.

Citation: AI in human–computer gaming: Techniques, challenges and opportunities (2023, December 29) retrieved 9 May 2024 from <u>https://techxplore.com/news/2023-12-ai-humancomputer-gaming-techniques-opportunities.html</u>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.