Proceedings of CONF-CDS 2026 Symposium: Data-Centric AI Security: Securing Models, Learning Agents, and Autonomous Systems
DOI: 10.54254/2755-2721/2026.CH33358

A Review of Artificial Intelligence in Fighting Games: Technical Evolution and Commercial Applications

Yantong Chen

NOIC Academy, Markham, Canada
chenyt0716@outlook.com

Abstract. Fighting games serves as an ideal testbed for real-time decision making AI due to the strict time restrictions and imperfect-information nature. Addressing the core challenges such as limited response time and restricted computational resources has significant value for real-time applications like autonomous driving. This paper provides a systematical review of the characteristics and evolutionary trajectory of AI in fighting games. It analyzes the transformation from early rule-based systems to modern Deep Reinforcement Learning (DRL) approaches, evaluating the respective advantages, limitations and applicable scenarios. A key finding is the successful application and prospect of DRL, exemplified by the Heterogeneous Exploitation Self Play (HESP) in the commercial fighting game Naruto Mobile,which overcomes the traditional scalability and generalization issues, making a significant advancement for AI in complex and real-time environments. Through analysis, DRL demonstrates potential in solving real-time decision-making issues and becomes a new direction for researchers.

Keywords: AI, fighting games, real-time decision-making, deep reinforcement learning (DRL)

1. Introduction

As a representation of a real-time two-player zero-sum game, fighting games possess the distinct research characteristics of real-time confrontation and rapid decision response, which have become an important and ideal test platform for Artificial Intelligence (AI) [1]. Either in the competition of a fighting game hosted by the IEEE Conference on Computational Intelligence Games or the FightingICE test platform, the strict limitation of the response time of actions of AI is less than or equal to 16.6 ms or 16.67 ms [2,3]. Meanwhile, based on the transitory response time, several core challenges of fighting games include the necessity for real-time decision-making under strict time constraints, the complexity of acting with incomplete information due to the simultaneous player moves, and the vast, dynamic state-action space unique to each character's move set [1,4]. In the context of limited computational resources, the high-intensity decision environment provides various algorithms and frameworks with a valuable test set for evaluating algorithmic response speed and assessing adaptive capacity.

In recent years, fighting games have attracted increasing attention from researchers in AI and game design, and the development of AI for fighting games has made significant progress. From the

DOI: 10.54254/2755-2721/2026.CH33358

early Finite State Machine (FSM) to the current large-scale application of Deep Reinforcement Learning (DRL), AI in fighting games has undergone a series of technical transformations. Unlike AI in Go or card games, which rely on massive computational resources and a turn-based decision-making mode, the real-time requirements of fighting games are more in line with the demands of real-time decision-making applications in reality, such as the reaction speed and capacity of autonomous vehicles in emergent situations. In addition, research in this field is diverse and lacks systematic reviews.

Thus, this paper aims to provide a systematic analysis of current conditions and the evolutionary trajectory of AI technology in fighting games, offering references and inspiration for researchers in related fields. The following sections focus on the core challenges, the trajectory of AI in fighting games, traditional example analysis, and the conclusion.

2. Core challenges

Due to the special real-time gameplay, sophisticated combo attacks, and intensive psychological gameplay, fighting games face a series of distinctive challenges that have become the dominant issues and focus of the genre. Resolving issues is important in fighting games and also provides reference value for other real-time decision-making fields.

To test AI algorithms and frameworks and improve their performance against core challenges, researchers and developers frequently use FightingICE as a test platform, which was chosen for the Fighting Game AI Competitions (FTGAIC) series from 2015 [3] and has been used hitherto. Originally introduced in 2013 [5] as an AI testbed specifically for fighting games, FightingICE allows researchers and students to test and evaluate their AI algorithms and frameworks [6]. After a series of improvements, the platform underwent a significant upgrade in 2022, becoming a sound-design-enhanced version: DareFightingICE. The new version adds a track for sound design [7] but retains the track for traditional fighting game AI.

2.1. Real-time decision-making and limited response time

FightingICE asks the AI to check 56 actions and choose one in 16.67 milliseconds for each frame. It sets a 15-frame delay (about 250 milliseconds) to copy the reaction speed of real human players [3]. So this makes the difficulty of making decisions even higher. Under strict time limits, AI programs need to finish all steps of knowing the game state, making decisions and carrying out actions. This requires the algorithms to work in a very efficient way.

To deal with this problem, earlier researchers put together Monte Carlo Tree Search (MCTS) and Evolutionary Algorithm (EA). They made a mixed fighting game AI that uses the strong points of both algorithms to improve the AI's reaction speed and ability [2,4]. But a new algorithm made by Tang and his teammates came out in 2020. The Rolling Horizon Evolution Algorithm Opponent Model (RHEAOM) is made better with the Rolling Horizon Evolution Algorithm (RHEA). The fundamental framework of RHEA is the optimization process of action sequence in limited time domain through a thorough evolution with selection, crossover and mutation [3], then select the first action with the highest fitness in the candidate action sequences [8,9], balancing the demand between real-time decision-making and long-term plan. The AI based on RHEA already has an average win rate of over 60% against AI bots in the 2018 FTGAIC of the IEEE Conference on Games, using three specific characters as experimental subjects (GAR, LUD, ZEN) [3].

DOI: 10.54254/2755-2721/2026.CH33358

2.2. Imperfect-information game

Fighting games feature elements of an imperfect-information game. AI cannot accurately capture the opponents' upcoming actions. The asymmetry of information influences the judgment of AI to a large extent. Given the existence of an imperfect-information game, the Nash equilibrium is a strategy for dealing with it. In an equilibrium state, no player can increase their payoff by unilaterally changing their strategy (assuming the strategies of other players remain unchanged) [10]. AI lacks the mental activities of humans and cannot gain significant advantages through psychological games like real human players do in fighting game competitions. Hence, developing an AI algorithm capable of opponent prediction and analysis, like real human players, is necessary for further progress in AI's real-time decision-making capability. Additionally, the indeterminacy and potential information correspond to the decision environment in reality.

To address issues with opponents, researchers use multi-objective evolutionary neural networks to enhance AI's generalization when facing diverse opponents, achieving a higher win rate in FTGAIC than traditional evolutionary neural networks. However, the win rate is only 1% against MCTS-based Thunder bots from the 2018 FTGAIC [11]. Thus, a new opponent model, based on a neural network and combined with RHEA, emerged in 2020 to predict the opponent's next action. The average win rate of RHEAOM-based AI is more than approximately 70% against 2018 FTGAIC bots. After introducing real-time learning algorithms, RHEAOM with policy gradient achieves the best performance among RHEAOMs in conjunction with three different real-time learning algorithms. The average win rate of RHEAOM with policy gradient is over 80% against the 2018 FTGAIC bots and over 60% against the opponent-model-introduced variants of MCTS-based Thunder bots [3]. This is a breakthrough in fighting game AI and represents significant progress in real-time decision-making, shifting attention from MCTS to RHEA.

2.3. Restriction of computational resources

Fighting games for commerce and the mass market are another direction for researchers to delve. Unlike a scientific research environment or a competition platform with sufficient computational resources, the operation of commercial fighting games on users' hardware platforms is constrained by limited random-access memory, limited computational resources, and other constraints. AI in commercial fighting games needs to run on consumer-grade hardware to provide a better player experience. This restricts the algorithmic complexity and computational intensity.

For this issue, game companies and researchers have proposed multiple optimization strategies. For instance, Tencent Games encountered the challenge of training AI due to the large number of characters with unique designs in Naruto Mobile (2016), and traditional DRL proved costly. To resolve the issue, Heterogeneous Exploitation Self-Play (HESP) was proposed, significantly reducing computational costs and increasing training efficiency by improving the generalization of the fighting game AI. It not only has the commercial value of player retention rate but is also the first time DRL has been applied in commercial fighting games [12].

3. Technical development of AI application in fighting games

Game AI shows smart behaviors by using algorithms from different fields. It gives players a very real and in-depth game experience [13]. The growth of AI in fighting games has three main stages. Each stage includes three kinds of systems. This shows a development path from rule-based to learning-driven AI. Learning about this technology change helps researchers and developers find the inner rules and future of AI development in fighting games.

3.1. Rule-based system

Early fighting games used AI that is based on FSM, Decision Trees and other algorithms, and these algorithms follow simple if-then rules set by developers beforehand to decide how the AI responds in certain game situations. This kind of AI has the benefits of simple programming, a fixed working system and high working speed, so developers can control AI moves and performance in a more accurate way. However, this AI also has clear shortcomings. AI built on FSM does not have the ability to adjust to changes, its reaction speed is much faster than that of human players, and its moves are easy to be predicted [6,14]. These unique shortcomings reduce the long-term fun and challenge for players, because the code needs a lot of manual changes and it is hard to fit different game environments and character settings.

This system is widely used in early fighting games with limited computing ability, such as Street Fighter 2 (1991), The King of Fighters 97 (1997) and other classic early fighting games. In modern times, this system is used to make training partners for special game guides, since AI moves are steady and easy to control, which is good for new players to learn basic operations and continuous attacks.

3.2. Search and planning-based system

As computing power becomes stronger, MCTS and other algorithms based on search and planning are used in fighting games. This system looks for possible decision paths by forward modeling and picks the best result from statistical data. Compared with rule-based systems, it is more flexible, adaptable and general [15]. MCTS can also work with different algorithms like Genetic Programming to get better results. But this system needs lots of forward statistical sampling and clear forward models. These things need strong computing resources and are hard to use in environments with huge state spaces and complex interactions [3,15].

This system is good for AI scenes with medium or high difficulty and quick tactical choices, such as hard levels or boss fights in Mortal Kombat 11 (2019) and Tekken 7 (2017). In game development, it can also simulate game plays, check character abilities and skill balance, and give useful data to help developers do their work.

3.3. DRL-based system

The appearance of DRL means that fighting games have come into a new stage. DRL brings together Deep Learning (DL) and Reinforcement Learning. DL deals with high-dimensional inputs, and RL makes decision strategies better. The DRL system can make its decision-making strategies better and better without much manual work. So it can fit in changing and unsure environments. But DRL also has weak points. This system gets better by trying and making mistakes. In real and complex uses, its strategy exploration is not efficient, and it needs a large number of samples. This makes training cost very high. DRL uses nonlinear work, so people cannot track its decision logic, and it is not easy to understand clearly [16].

The system is primarily used in modern fighting games or research platforms that require extremely high AI performance. Commercial fighting games with a large number of characters, like Naruto Mobile (2016), use DRL to significantly reduce computing costs and training time through

DOI: 10.54254/2755-2721/2026.CH33358

large-scale self-play training [12]. In FTGAIC, DRL is also used to improve AI's performance in the battle when confronting a new opponent. For instance, exploiting policy gradient in RHEAOM enhances competitiveness by predicting the opponent's actions in competitions [3].

As shown in the table 1, the three categories are the rule-based system, the search-and-planning-based system, and the DRL-based system. Each category has representative technologies, advantages, limitations, and applicable scenarios.

Table 1. The comparison of AI systems in fighting games

Category	Representative technologies	Advantages	Limitations	Applicable scenarios
Rule-based System	Finite State Machine, Decision Tree	easy programming, deterministic system, high efficiency	Adaptability scarcity, overquick reaction time, predictable	Early fighting games, training partner
Search and Planning-based System	MCTS	flexibility, high adaptability, good generalization ability	Large-scale sampling, high computational resources limited adaptation	Difficult levels and boss battles, test tool in development
DRL-based System	DRL	no required manual intervention, relatively strong adaptation	low exploration efficiency, large required samples, poor interpretation	Commercial fighting games, complex real-time decision environments

4. Application: DRL in commercial fighting games

With the rapid advances in algorithms and computational capacity, DRL is gradually moving from academic research into the development of commercial fighting games, bringing essential reform to the genre. In this process, Naruto Mobile (2016) by Tencent Games is the first commercial fighting game to apply DRL [12]. Not only has the feasibility of this technology been verified, but it has also established a leading position.

Based on the Japanese anime Naruto, Naruto Mobile is a one-on-one mobile fighting game published by Tencent Games, with more than 100 million players registered. AI is widely used in game scenes such as level challenges and player training, and it needs to copy different fighting styles for players at different levels to keep game difficulty in a balanced state. However, traditional DRL has a big problem of needing huge amounts of training in this kind of game. Different from traditional fighting games that have about 20 characters, Naruto Mobile has more than 300 different characters, and each character has special attack moves and unique skills. The complex real-time decision-making process and the huge strategy space both make computational cost and training time rise sharply. The most important goal of designing a large-scale DRL training method is to solve the scalability problem that exists in traditional DRL [12].

Heterogeneous Exploitation Self-Play (HESP) was created by researchers to solve this exact problem. It is a special DRL method that is made for large-scale training, and its core goal is to make the generalization ability of AI agents stronger and make training efficiency better. In the training process, HESP solves the weak points of traditional methods by training two kinds of agents together, which are ID-based agents and quantitative-based agents. ID-based agents depend on identification information like character IDs and skill IDs to explore the features of different

DOI: 10.54254/2755-2721/2026.CH33358

characters and quickly learn the best strategies for certain characters. Quantitative-based agents cut down their dependence on IDs and realize cross-character generalization, and this lets AI fit in with unknown characters by turning states into numerical features, such as character attributes, hitbox, skill cooldown times and other features [12]. The joint work of these two kinds of agents makes sure that AI performs well when facing known characters and also lets AI fit in with unknown characters very quickly.

The whole training process is divided into two different stages, which are the independent training stage and the hybrid training stage. In the first stage, the self-play method is used to train ID-based agents and quantitative-based agents separately. ID-based agents use the ID of each character to use different strategies, and they quickly raise their win rate when fighting against known characters. Quantitative-based agents explore common rules and patterns by using numerical information. Besides, this stage uses the Importance Weighted Actor-Learner Architecture [17] algorithm to offer various experience sources for agent learning. In this stage, agents need to do many iterations, and each iteration is compared with the last one until the win rate reaches 90%. In the second stage, because the training efficiency of quantitative-based agents is lower than that of ID-based agents, hybrid training is used to improve the efficiency of quantitative-based agents and raise their win rate when fighting against unknown characters. The best policy of ID-based agents is added to the policy pool of quantitative-based agents, which offers high-quality opponent samples and helps quantitative-based agents fit in with difficult opponents quickly and spread their skills to other characters [12]. Generally speaking, ID-based agents solve the problem of targeted learning efficiency, while quantitative-based agents solve the problem of generalization ability. Working together, they help HESP make big breakthroughs in AI training for large-scale fighting games, and make sure AI has high performance for known characters and can fit in with unknown characters quickly.

Compared with the traditional ID-based method, the win rate of this method against unknown opponents reaches 66%, which is 15% higher than before. Its generalization ability is improved a lot. In addition, the training time for each iteration is cut down to 4 hours, and it supports the synchronous training of more than 300 characters, so it solves the problem of high training costs. The 5-day retention rate of new players has gone up by 15%, and the 30-day retention rate has increased by 4% [12]. HESP is the first method that applies DRL to large-scale commercial fighting games successfully. It not only keeps the targeted feature of ID-based methods but also uses the generalization advantages of quantitative-based agents, and offers an expandable model for AI training in complex decision-making scenes.

5. Conclusion

This paper systematically reviews the development trajectory, core challenges, and application of artificial intelligence technology in fighting games. As an ideal test platform for real-time decision-making, the research on fighting game AI focuses on the core issues such as strict response time limits, imperfect-information games, and limited computing resources. The solutions to these problems have significant reference value for real-time decision-making fields in reality, such as autonomous driving.

The evolution of technology presents a clear path from rule-based systems to search and planning-based systems, and ultimately to DRL systems. Among them, DRL, with its strong environmental adaptability, is moving from academic research to commercial applications. Taking the HESP method, which was first applied on a commercial fighting game, Naruto Mobile, as an example, it effectively addresses the high training costs and insufficient generalization ability of

DOI: 10.54254/2755-2721/2026.CH33358

traditional DRL in commercial games with a large number of characters and diverse skills by collaboratively training agents based on IDs and quantitative features, significantly improving training efficiency and the AI's adaptability to new characters.

In conclusion, the research on AI in fighting games not only pushes the performance AI agents but also provides a reference paradigm for a wider range of real-time decision-making systems through algorithmic innovations in real-time, uncertainty, and resource-constrained environments. In the future, further improving model efficiency, interpretability, and generalization capabilities in more complex and dynamic environments will be the continuous exploration directions in this field.

References

[1] Z. Tang, R. Liang, Y. Zhu, and D. Zhao, "Intelligence decision making approaches for real time fighting games, " Control Theory & Applications, vol. 39, no. 6, pp. 969-985, 2022.
[2] M. J. Kim, and C. W. Ahn, "Hybrid fighting game AI using a genetic algorithm and Monte Carlo tree search, " in Proceedings of the Genetic Evolving Computational Conference Companion (GECCO), 2018, pp. 129-130.
[3] Z. Tang, Y. Zhu, D. Zhao, and S. M. Lucas, "Enhanced rolling horizon evolution algorithm with opponent model learning, " IEEE Transactions on Games, 2020.
[4] G. T. Lam, D. Logofatu and C. Badica, "A novel real-time design for fighting game AI, " Evolving Systems, vol. 12, no. 1, pp. 169-176, 2021.
[5] F. Lu, K. Yamamoto, L. H. Nomura, S. Mizuno, Y. Lee, and R. Thawonmas, "Fighting game artificial intelligence competition platform, " in 2013 IEEE 2nd Global Conference on Consumer Electronics, 2013, pp. 320–323.
[6] G. Martinez-Arellano, R. Cant, and D. Woods, "Creating AI characters for fighting games using genetic programming, " IEEE Transactions on Computational Intelligence and AI in Games, vol. 9, no. 4, pp. 423-434, 2016.
[7] I. Khan, T. V. Nguyen, X. Dai and R. Thawonmas, "DareFightingICE Competition: A Fighting Game Sound Design and AI Competition, " in 2022 IEEE Conference on Games (CoG), 2022, pp. 478-485.
[8] D. Perez, S. Samothrakis, S. Lucas, and P. Rohlshagen, "Rolling horizon evolution versus tree search for navigation in single-player real-time games, " in Proceedings 15th Annual Conference Genetic Evolutionary Computation, 2013, pp. 351-358.
[9] R. D. Gaina, S. M. Lucas, and D. Pérez-Liébana. "Population seeding techniques for rolling horizon evolution in general video game playing, " in 2017 IEEE Congress on Evolutionary Computation (CEC), 2017, pp. 1956-1963.
[10] D. M. Kreps, "Nash equilibrium, " The new Palgrave dictionary of economics. Palgrave Macmillan, 2018, pp. 9251-9258.
[11] S. Künzel, and S. Meyer-Nieberg, "Coping with opponents: multi-objective evolutionary neural networks for fighting games, " Neural Computing and Applications, vol. 32, no. 17, pp. 13885-13916, 2020.
[12] E. S. Liu, W. Li, Y. Zhou, H. Cao, and Z. Zeng "Naruto mobile: Ai sparring partner using heterogeneous deep reinforcement learning, " in 2023 IEEE Conference on Games (CoG), 2023, pp. 1-8.
[13] A. Filipović, "The role of artificial intelligence in video game development, " Kultura polisa, vol. 20, no. 3, pp. 50-67, 2023.
[14] C. Fairclough, M. Fagan, B. M. Namee, and P. Cunningham, "Research directions for ai in computer games, " in Proceedings 12th Irish Conference Artificial Intelligence Cognitive Science, 2001, pp. 333–344.
[15] M. Świechowski, K. Godlewski, B. Sawicki, and J. Mańdziuk, "Monte Carlo tree search: A review of recent modifications and applications, " Artificial Intelligence Review, vol. 56 no. 3, pp. 2497-2562, 2023.
[16] Y. Qiu, "Deep reinforcement learning research and applications, " ITM Web of Conferences, vol. 78, Art. no. 01010, 2025.
[17] L. Espeholt, H. Soyer, R. Munos, K. Simonyan, V. Mnih, T. Ward, Y. Doron, V. Firoiu, T. Harley, I. Dunning, S. Legg, and K. Kavukcuoglu, "IMPALA: Scalable distributed deep-RL with importance weighted actor-learner architectures, " in Proceedings of the 35th International Conference on Machine Learning, ser. Proceedings of Machine Learning Research, J. Dy and A. Krause, Eds., vol. 80. PMLR, 10–15 Jul 2018, pp. 1407–1416.