Reinforcement Learning in Digital Games: An Exploration of AI in Gaming

Ziwei Tang

School of Advanced Technology, Xi'an Jiaotong Liverpool University, Xi'an, 710000, China
Email: Ziwei.Tang18@student.xjtlu.edu.cn

Abstract. Reinforcement Learning (RL), a pivotal technique in AI, finds extensive applications in game AI, ranging from elementary board games to sophisticated strategy games. This application not only carries significant practical implications but also catalyzes theoretical advancements in AI. The objective of this paper is to conduct a comprehensive review and analysis of the current application of RL in game AI, with the intent of uncovering its potential for augmenting the gaming experience. This research focuses on exploring five primary RL methodologies and their implementation in specific games. The aim is to gain a deeper understanding of the advantages, limitations, and impact of these methods on the gaming experience. The insights gleaned from this research endeavor will foster a more nuanced understanding of the practical implementation of RL, serving as an invaluable resource for subsequent research and development pursuits. Additionally, this paper analyzes and discusses the challenges and characteristics that AI games may face in the future. Finally, the paper concludes with a summary and prospects for future research.

Keywords: Reinforcement learning; digital games; AI

1 Introduction

As the transition through the 21st century progresses, the pervasive influence of AI is palpable in all spheres of life, most notably in the domain of digital gaming. RL, a foundational component of AI, demonstrates expanding applicability in game AI across various games of differing complexity. RL's role is dual-fold—it holds immense practical relevance and concurrently fuels AI's theoretical evolution. Observing AI's superior performance in board games instigates a paradigm shift (as depicted in Fig. 1. The impressive performance of AI not only substantiates its formidable capabilities, but also provokes reflections on harnessing AI to improve human abilities. This presents avenues for reciprocal advancement in the landscape of human-AI interaction.

P. Kar et al. (eds.), *Proceedings of the 2023 International Conference on Image, Algorithms and Artificial Intelligence (ICIAAI 2023)*, Advances in Computer Science Research 108,
https://doi.org/10.2991/978-94-6463-300-9_37

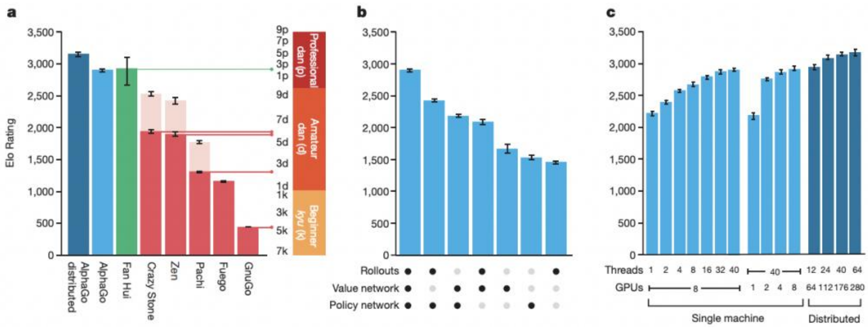
Fig. 1. Tournament evaluation of AlphaGo [1].

RL, as a subset of machine learning (ML), empowers an AI system to determine the optimal course of action within a specified context through exploration and experimentation within the environment. This approach has substantial potential applications, one of which is the evolution of more intelligent and lifelike game AI. Despite the existence of seminal research within this sphere, a significant amount of controversy and uncertainty remains regarding the best application of RL techniques for game AI enhancement [2][3].
Consequently, the objective of this paper is to conduct a comprehensive review and analysis of the current application of RL in game AI, with the intent of uncovering its potential for augmenting the gaming experience. The exploration will focus on the five primary RL methodologies and their implementation in specific games. The aim is to achieve a deeper understanding of the advantages and limitations of these approaches, as well as their influence on the gaming experience. The insights from this research will be valuable for future research and development pursuits [4].

2 Key Algorithms in Game AI: A Comprehensive Analysis

The application of various algorithms and ML techniques in game AI is reflected primarily in determining the optimal solutions for future states of Non-Player Characters (NPCs) or in-game objects, and their consequent actions.

2.1 Search Algorithm

The primary use of search algorithms in game AI is for discerning the best possible solution for the future state of NPCs or in-game objects. For example, the A* algorithm functions as a heuristic search method for identifying the most optimal and shortest path from the start to the end point by accumulating and evaluating potential path alternatives. This strategy greatly reduces computational time and enhances efficiency, enabling NPCs to swiftly adapt to changing game situations.

2.2 Planning Algorithm

Planning algorithms furnish NPCs with various behavior planning schemes based on a pre-determined strategy, enabling game characters to decide based on the current environment and goals, and act accordingly. One widely used planning algorithm is the STRIPS planning algorithm, which plays a significant role in intelligent decision-making by game AI. In strategy games, the algorithm can help the AI anticipate potential enemy behavior and devise preemptive counterstrategies, thus enhancing the game's challenge.

2.3 Reinforcement Learning

RL is predicated on forward-looking objectives, leveraging trial-and-error tactics to stimulate AI learning and optimization. RL utilizes environmental feedback and modifies behavior in response to rewards or punishments. A notable example of RL is the Monte Carlo Tree Search (MCTS), increasingly applied in numerous gaming areas due to its stellar performance in board games such as Go. RL bolsters the autonomous learning and adaptability of AI, making a pronounced impact on the long-term challenge of games. The DeepMind team's research provided a tangible illustration of RL's effective application in a real game environment, where the team successfully trained an AI to perform at a high level in the StarCraft II game (see Fig. 2). This instance robustly affirms the broad applicability and efficacy of RL techniques in complex game environments.

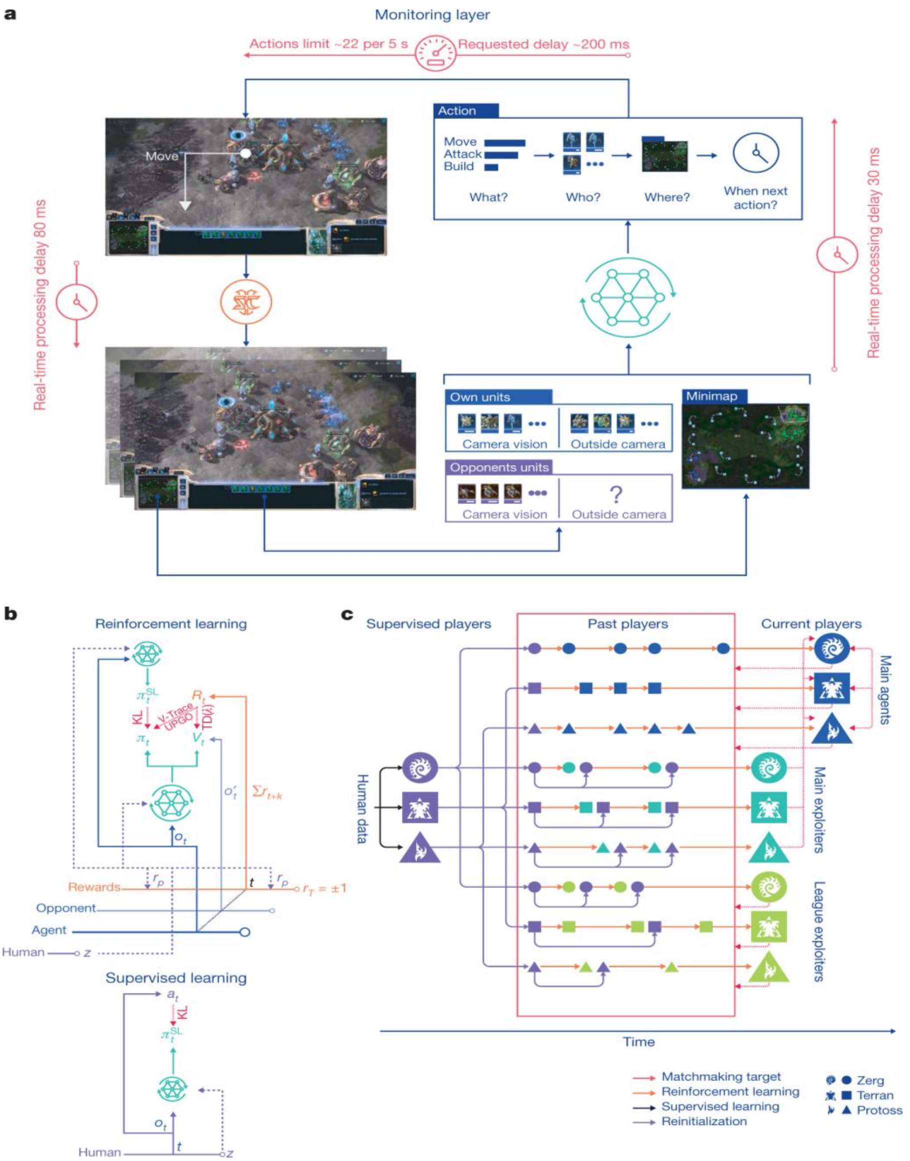

Fig. 2. AlphaStar looks at the game through an overview map and a list of units. Its action output includes the type of action to be issued, the application object, the target location, and when the next action will be issued. AlphaStar sends actions to the game through a monitoring layer to limit the frequency of actions [5]

2.4 Finite State Machine

The Finite State Machine (FSM) is a behavioral model premised on state transitions, extensively employed in game AI design. The FSM facilitates the behavior scheduling and state transitions of NPCs by systematically illustrating all potential states and their transition relationships. FSM, as a simple and efficient method, provides limited and predictable outcomes, ensuring the stable functioning of the game. However, FSM might encounter limitations in more intricate game worlds.

2.5 Behavior Tree

The behavior tree is a tree-based behavior control methodology that organizes game character behaviors in a hierarchical manner. Its key components encompass elements such as condition, behavior, sequence, and selection. Compared to FSM, behavior trees exhibit substantial flexibility in complex logic and multi-conditional tasks. By dynamically allocating task priorities, the behavior tree affords NPCs enhanced adaptability and intelligence, thereby augmenting the sense of realism and immersion in the game.

The success of AlphaGo in using the MCTS algorithm for decision-making underscores the critical role of search algorithms in gaming. This algorithm is particularly adept at analyzing and planning the variety of game states that may transpire in a board game (see Fig. 3). This discovery that has profound implications for our a profound understanding and usage of search algorithms' potential in game design.

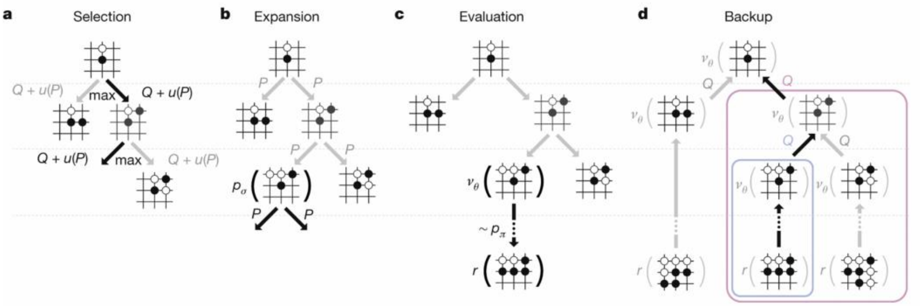

Fig. 3. MCTS in AlphaGo [1]

This section elucidates the distinct role that various methods play within the realm of gaming, highlighting their transformative impact on the presentation and experiential dimensions of contemporary games. It is imperative to acknowledge that method selection is commonly contingent on the game type and design objectives, to yield the most suitable amalgamation of outcomes. Furthermore, in the course of practical implementation, these methods are frequently amalgamated to foster an elevated level of game AI design and performance optimization.

3 Exploring AI Application in Contemporary Games

In this section, this paper examines the AI technologies employed in five renowned games, and the transformative effect they exert on diverse gaming genres. This discussion revolves around their applications in the action game "Super Mario Bros", the strategy game "StarCraft II", the role-playing game "The Elder Scrolls V: Skyrim", the simulation game "Cities: Skylines", and the sports game "FIFA".

- Action Game - "Super Mario Bros":

AI technology applications in this iconic game encompass search algorithms, planning algorithms, and rule-based methods. These methodologies not only escalate the game's challenge but also optimize the game design process. A prime example is the automatic generation of level solutions, which enhances the game's playability by alleviating the designer's workload [6]. Concurrently, the implementation of search and planning algorithms has sparked innovative designs for adversaries and impediments, affording players a more challenging and enjoyable gaming experience [7].

- Strategy Game - "StarCraft II":

In this warfare strategy game, the AI incorporates techniques like deep learning, RL, and Monte Carlo tree searching to tackle the game's complexity [8]. These techniques equip AI players with superior strategic capability and swift response times, making in-game combat extremely challenging. Furthermore, with extensive training, the AI can manage a broad spectrum of strategic situations, thereby offering a more diversified and intriguing gaming experience for human players.

- Role-playing Game - "The Elder Scrolls V: Skyrim":

This open-world game teeming with adventure uses AI for NPCs through methods like FSM, goal-oriented action planning (GOAP), and behavior trees. The integration of these methods into games can render NPCs more realistic and dynamic, enabling them to adapt their behavior to different situations. This not only intensifies the game's immersion and player engagement but also furnishes a valuable benchmark for world-building and narrative experience.

- Simulation game - "Cities: Skylines":

AI technology applications in this urban development simulation game comprise rule-based systems, procedural generation, and agent control. By auto-adjusting in-game factors like building, traffic, etc., the game's complexity and challenge are significantly heightened. Consequently, players are compelled to invest more time and effort into resolving AI-generated problems, culminating in an enhanced gaming experience.

- Sports Game - "FIFA":

In the globally acclaimed football game "FIFA", noteworthy strides have been made in player behavior and team collaboration. The AI employs techniques like FSM, coordinated control, and hierarchical learning to facilitate systematic communication and coordination among players. These methods allow NPCs to express and

collaborate more naturally, making the game feel more realistic and challenging, and thereby enticing players to dedicate more time to gameplay.
In conclusion, AI technology's wide application across varied game types positively influences the challenge, realism, immersion, and efficiency of game design. Studying these successful cases is anticipated to provide valuable insights for the future evolution and innovation of AI in the gaming industry. However, it is critical to note that during the application of AI technology, different game types' characteristics and needs should be fully considered in the design process to ensure an optimal gaming experience.

4 The challenges of AI games

The incorporation of specific methodologies in game AI design poses a plethora of challenges. This section seeks to dissect the conundrums associated with the application of game AI, focusing on the dimensions of AI design scope, game types, and player requirements.

- Diversity Demands:

The diversity inherent in the modern computer game landscape necessitates significantly variant game AI needs. These range from simple puzzle games to intricate strategy games, each of which has vastly differing game AI design objectives. For instance, while puzzle games concentrate on problem-solving and algorithm optimization, role-playing games place an emphasis on character personalities and story arcs. Therefore, the selection of an appropriate method should be flexible and customized to cater to the specific game genre and player requirements.

- Computational Resource Constraints:

Computational resources emerge as a critical factor in game AI design. Methods such as search algorithms, planning algorithms, and RL often demand extensive computation. However, the real-time requirements of the gaming environment and hardware limitations necessitate an AI approach that strikes a balance between efficiency and performance. Consequently, achieving high AI performance under the constraint of limited resources presents a substantial challenge.

- Complexity and Scalability:

The continuous evolution of the gaming industry implies that the game world's scale is expanding, thereby escalating player demands. Designing scalable game AI for these complex environments remains a challenge. When confronting complex situations, FSM might fail to achieve high intelligence. While behavior trees hold an advantage in extensibility, their design process necessitates the articulation of specific logic, which can lead to a design burden. The optimal method selection and its practical application will hinge on the game context and its core objectives. In addressing player needs, it's imperative not just to focus on their gaming skills or play duration, but also to consider their emotional experience. The emotional game loop model illustrates how player's emotional feedback can be employed to generate game content, which is vital for comprehending and designing more human-centric game experiences.

- Balancing Player Experience and AI Intelligence:

Different players have diverse preferences and needs. Therefore, game AI must strike the right balance to challenge the player without rendering the game overly difficult or monotonous. Additionally, excessive intelligence can lead to AI over-optimization, potentially diminishing the gaming experience or making the game less entertaining. Hence, the equilibrium between player experience and AI intelligence needs thorough consideration in the practical application of the method.
In conclusion, the application of specific methods in game AI design poses numerous challenges. To address these, researchers must persistently explore more optimized and efficient algorithms and design methodologies to develop continually evolving game AI systems capable of meeting escalating player demands.

5 Artificial Intelligence in Gaming: Critical Attributes

The focus of this section is to delineate the significant traits necessitated by the incorporation of AI in gaming. This discussion will encompass six crucial characteristics.

As the gaming industry experiences robust growth, the incorporation of Artificial Intelligence (AI) in this realm has attained mounting significance. The essence of AI within games is to forge diverse, engaging, and challenging gameplay experiences. A survey of relevant literature allows us to crystallize six salient characteristics of game AI:

1. Adaptability: Of paramount importance to game AI is adaptability. AI should possess the capacity to discern changes in the gaming environment and player behaviors, subsequently formulating apt decisions or altering strategies. In strategy games, for example, AI should devise diverse tactical deployments, contingent on terrain, resources, and war conditions.

2. Autonomy: Game AI necessitates the ability to independently arrive at decisions and undertake actions based on inbuilt rules and procedures, irrespective of player inputs. Often, game AI adopts a behavior tree or state machine architecture to guarantee autonomous operation in myriad situations.

3. Interactivity: Game AI should be capable of engaging meaningfully with other AI units and players. This involves effective collaboration in multiplayer games and appropriate responses to player behavior [9-10].

4. Intelligence: Intelligence emerges as a fundamental attribute of game AI. AI should display the capacity to learn, optimize strategies, and behavior through self-learning mechanisms. Technologies such as Deep Learning and RL have positively contributed to enhancing game AI intelligence [9].

5. Simulation: Game AI is conceived to mimic real-world behavior, inclusive of compliance to rules and real-world physical consequences. To facilitate this, game AI typically employs simulation systems encompassing physical, psychological, and ecological models.

6. Virtuality: Game AI, a virtual entity, derives its behavior from procedural logic. This virtuality might create a chasm between the behavior of game AI and real-world actions. However, through ceaseless refinement and enhancement, the virtual behavior of game AI can increasingly approximate real-world behavior.

In conclusion, within the gaming universe, AI design should judiciously take into account the attributes of adaptability, autonomy, interactivity, intelligence, simulation, and virtuality to engineer an increasingly realistic, captivating, and challenging gaming experience.

6 Conclusion

In recent years, reinforcement learning has made significant advancements in the domain of game AI. This article provides an overview of five prominent reinforcement learning methodologies, coupled with detailed case studies on specific games. The progress in these reinforcement learning techniques, while substantial, merely scratches the surface of their potential impact. It is evident that, while individual reinforcement learning methodologies offer unique advantages in certain game environments, the future might lean towards hybrid models, amalgamating the strengths of these disparate methods to fashion more responsive, intelligent, and immersive gaming ambiances.

Despite the progress to date, fully harnessing the potential of RL remains an elusive goal. Especially when confronted with intricate game scenarios, multiplayer interactions, and expansive open worlds, further research and exploration are imperative. Anticipated future studies will likely focus on these arenas, deepening the understanding of the interplay between reinforcement learning and artificial intelligence within games.

However, this trajectory is not devoid of challenges. With the continuous evolution of digital games, particularly considering the advent of virtual reality and augmented reality platforms, the requisites for game AI will burgeon in complexity. Games are no longer static entities; they evolve contingent upon player behaviors, global gaming communities, and real-world incidents. This dynamic nature necessitates AI systems endowed with both adaptability and robustness, poised to tackle a spectrum of challenges head-on.

In summation, this article aspires to furnish a comprehensive perspective for individuals intent on applying reinforcement learning to game AI. While our grasp of this domain has ripened, its horizon brims with both challenges and prospects. We advocate for researchers and developers to invest amplified efforts in this sphere, such as exploring population-based deep reinforcement learning in first-person multiplayer games, for it is through collaboration that we can truly unlock reinforcement learning's potential to amplify game AI and deliver an enhanced gaming experience.

References

1. Silver, D., Huang, A., Maddison, C. J., Dieleman, S. (2016). Mastering the game of Go with deep neural networks and tree search. Nature, 529(7587), 484-489.
2. Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, (2018). A general reinforcement learning algorithm that master's chess, shogi, and Go through self-play. Science, 362(6419), 1140-1144.
3. Arulkumaran, K., Deisenroth, M. P., Brundage, M., & Bharath, A. A. (2017). Deep reinforcement learning: A brief survey. IEEE Signal Processing Magazine, 34(6), 26-38.
4. Li, Y. (2017). Deep reinforcement learning: An overview. arXiv preprint arXiv:1701.07274.
5. Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, (2019). Grandmaster level in StarCraft II using multi-agent reinforcement learning. Nature, 575(7782), 350-354.
6. Bontrager, P., Khalifa, A., Mendes, A., & Togelius, J. (2018). Automatic playtesting for game parameter tuning via deep reinforcement learning. arXiv preprint ArXiv: 1802.06881.
7. Polydoros, A. S., Nalpantidis, L., & Kruger, V. (2017). Real-time reinforcement learning for AGV job-shop scheduling. The International Journal of Robotics Research, 36(2), 231-248.
8. Orkin, J. (2018). Three States and a Plan: The AI of FEAR. In Game AI Pro, CRC Press. pp. 159-180.
9. Jaderberg, M., Mnih, V., Czarnecki, W. M., et al. (2019). Human-level performance in first-person multiplayer games with population-based deep reinforcement learning. arXiv preprint arXiv:1907.04543.
10. Badia, A. P., Piot, B., Kapturowski, S., et al. (2020). Agent57: Outperforming the Atari human benchmark. arXiv preprint arXiv:2003.13350.