Published by :
https://www.ijert.org/
An International Peer-Reviewed Journal

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 03 , March - 2026

Comparative Study Adaptive Difficulty Adjustment in Games using Reinforcement Learning: Enhancing Player Engagement Through Personalized Challenges

Armaan Jitesh Kumar Jain
School of CS & IT
Student
Jain Deemed to Be University
Bangalore.

Jeet Bhindora
School of CS & IT
Student
Jain Deemed to Be University
Bangalore.

Tarun K
School of CS & IT
Student
Jain Deemed to Be University
Bangalore.

Abstract - The balance of player skills and game difficulties is a crucial factor in maintaining engaging gaming experiences in modern digital games. A constant level of difficulty is not sufficient in accommodating varying and developing player skills; as a result, players may become bored and frustrated and eventually disengage from the game. This paper is a comparative analysis of adaptive difficulty adjustment techniques using reinforcement learning as a tool in personalizing game difficulties in real-time. The framework uses a Markov Decision Process in modeling the game environment; a reinforcement learning agent is used in observing performance metrics such as accuracy rate, success rate, and response time in adapting parameters such as enemy power level, task difficulty, and resource availability in maintaining a state of flow in players by adapting game difficulties in accordance with individual skills. Based on recent research in dynamic difficulty adjustment techniques, player modeling, and game agents, the analysis evaluates the advantages and disadvantages of RL in accommodating game difficulties in accordance with individual skills in comparison with traditional rule-based systems. The results of the analysis highlight RL as a promising tool in accommodating player-centric game design in achieving more adaptive and personalized gaming experiences.

2. INTRODUCTION

In the rapidly changing world of digital entertainment, static gaming experiences are increasingly being replaced with personalized gaming experiences that are more player-centric. The traditional “one size fits all” approach to difficulty settings often fails to consider the various skill sets, learning rates, play styles, and preferences of the players, resulting in an unenthusiastic experience with avoidable frustration. As the complexity of the gaming experience increases with the diversity of the gaming community, static difficulty settings chosen at the start of the game often do not remain relevant throughout the gaming experience. This often results in the gaming experience becoming too easy and boring, as well as too hard and frustrating, resulting in an unenthusiastic experience.

In this sense, the advent of advanced game analytics, player modeling, and artificial intelligence enables the analysis of player behavior and the adaptation of the difficulty level. Accordingly, the application of AI-based adaptive difficulty is centered on the analysis of the behavior of players with different skill levels, reaction rates, and strategic approaches to the challenges posed by the game. It is in this sense that reinforcement learning-based approaches can monitor the performance of players, including factors such as accuracy, success rates, and response rates, to adapt the enemy level, task level, and resource availability to individual players. Ultimately, the aim is to maintain the player in the optimal flow state.

However, there are also significant technical and practical concerns regarding the design and use of such adaptive difficulty control systems. Not only is it important to understand how to effectively represent the state of a player in a form that is both informative and tractable from a computational standpoint, but it is also important to understand how to design the reward function so that it is balanced and fair and how to prevent agents from learning and exploiting any unintended dynamics of a game. There are also important concerns regarding the overall transparency and fairness of such adaptive difficulty control systems and how they are perceived by the player.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 03 , March - 2026

The main focus of this research is to gain a better understanding of how reinforcement learning could be used to move away from static and rule-based difficulty management systems and towards adaptive systems in which game difficulty is constantly being selected and managed based on the player's performance and behavior. Through an analysis of different AI methods for managing difficulty and a focus on RL-based methods, the research hopes to illustrate the benefits of adaptive systems for players while also adhering to design and ethical constraints.

Research Objectives
Adaptive Difficulty Assessment and Control
Investigate the ways in which AI-based models can assess player performance in real time, categorize players according to their level of engagement or skill, and utilize such assessments to inform real-time adjustments to difficulty.
Reinforcement Learning for Personalized Challenge
Investigate the ways in which reinforcement learning can be used to represent the game as a decision process, where player performance metrics and state variables are used to make decisions regarding difficulty that maintain an optimal level of challenge for the player.

Optimization of Flow-Oriented Gameplay
Explore the possibility of designing adaptive systems to ensure the player remains in the flow state by adjusting factors such as enemy difficulty, task difficulty, and resource availability based on the data collected during the gameplay.

Practical, Ethical, and Design Considerations
Explore the major challenges facing the application of RL-based adaptive difficulty adjustment in actual games, including the computational cost, stability, and transparency of the algorithm, the level of trust the player should be encouraged to maintain, and the potential for manipulative behavior.

3. LIMITATIONS OF THE STUDY

Computational complexity and performance overhead
Reinforcement learning algorithms typically demand significant computational power, both for training and in real-time, while adapting to changing situations. In games, where frame rate is critical in real-time games, RL-based controllers may face challenges in meeting these demands, especially on platforms with limited capabilities, such as mobile devices.

Reward design and stability of learning the success of RL-based difficulty adjustment largely depends on a well-crafted reward system, which, if not designed properly, may result in unstable learning, where difficulty levels may oscillate, or a situation where the agent may become proficient in a particular aspect of the game, which, although beneficial in terms of statistics, may not provide a good gaming experience for the player.

Lack of transparency and explainability
One issue with RL is that the policy is generally learned rather than scripted, which means that the reasons for the changing difficulty are not immediately obvious to the designer, the tester, or the player.

Risk of overfitting to certain types of players
If the RL policy is optimized for certain types of players, then it is possible that the policy will work well for those players but not for players who do not fit the typical profile. This is not what we want with personalization, as we do not want to create an adverse experience for certain players.

Data requirements and the cold start problem
Adaptive difficulty requires some way to gauge the performance and interest of the player. In the absence of such information, the RL policy will not be able to make effective decisions, potentially resulting in an unstable difficulty level.
Unintentional manipulation and fairness
The direct intervention of RL-based systems in the challenges and progression of players may cause unintentional manipulation of players' behaviors and may not be fair to the players' interests. Ensuring the fairness of adaptive difficulty is a major problem.

Complexity of integration of RL with existing design workflows
Integrating RL into existing game development workflows is a complex process. The design process will have to be significantly changed to include RL-based systems and tools. The design process will have to take into account the changing nature of RL-based systems and the complexities involved in implementing such systems.

4. SCOPE OF THE STUDY

The specific topic of interest within this scope is adaptive difficulty adjustment in digital games, with a particular emphasis on reinforcement learning as a fundamental approach to game personalization. This paper will explore how reinforcement learning can be applied to model a game as a decision-making process, where various player performance metrics and game state variables inform real-time difficulty adjustments to maintain an optimal balance between challenge and skill. Within this scope, other popular AI techniques applied to game development, such as Finite State Machines, Behavior Trees, and Goal-Oriented Action Planning, will be considered primarily as a comparison baseline for reinforcement learning-based approaches.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 03 , March - 2026

The paper will be more conceptual and theoretical in nature, with a strong emphasis on existing research, frameworks, and applications, to explore how different AI techniques contribute to player-centric game design, particularly in terms of engagement, flow, and retention. This will be more focused on single-player or player-versus-environment game modes where game systems can adjust difficulty based on player behavior, as opposed to competitive multiplayer balancing scenarios.

The scope of this work is restricted to the design, capabilities, and challenges of reinforcement learning-based adaptive difficulty management. Topics such as state representation, actions, reward functions, and other related aspects of adaptive difficulty management are included in the scope. Other topics such as procedural content generation, story adaptation, monetization models are included only insofar as they are related to difficulty management. Various ethical considerations such as computational costs, unpredictability, transparency, and the possibility of manipulation or unfairness are mentioned on a general level in order to encourage responsible use of adaptive difficulty management, but legal considerations are not included in the scope.

5. LITERATURE REVIEW

Paraschos & Koulouriotis (2023) – "Game difficulty adaptation and experience personalization: A literature review"
Conduct a systematic review of 109 papers related to game difficulty adaptation and experience personalization (2005-2021) and categorize different techniques including DDA and PCG, and prove that adaptive systems can significantly enhance user experience in various domains including game environments, educational systems, and rehabilitation.

Zhang & Goh (2021) – "Personalized task difficulty adaptation based on reinforcement learning"
Present a framework based on RL for adapting task difficulty in MOOC environments using a bootstrapped policy gradient approach and a clustering-based difficulty ranking approach, which outperforms a traditional rule-based adaptation approach in a visual memory game.

Mortazavi, Moradi & Vahabie (2024) – "Dynamic difficulty adjustment approaches in video games: a systematic literature review"
Conduct a systematic review of different DDA techniques in video game environments, including machine learning techniques, player modeling techniques, and data sources including performance, emotion, and personality, and prove that DDA can significantly enhance enjoyment, flow, motivation, engagement, and immersion in game environments while highlighting the importance of understanding player preferences.

Lopes & Bidarra (2011) – "Adaptivity challenges in games and simulations: a survey"
To survey adaptivity in game and simulation environments in terms of purposes, targets, and methods of adaptivity, and highlight how content generation and semantic modeling can enable a move towards more dynamic and customized game worlds and quests.

Zohaib (2018) – "Dynamic difficulty adjustment (DDA) in computer games: A review"
To summarize DDA techniques in computer games in terms of adapting game features and scenarios in real-time based on the skill level of players in order to avoid boredom and frustration due to linear difficulty curves in game environments and highlight how DDA can act as a customized solution in game environments.

6. RESEARCH METHODOLOGY

This research is based on a literature-based research design, aiming to explore adaptive difficulty adjustment in digital games in general and reinforcement learning and associated AI techniques in particular.

The sources were identified based on a systematic search in major academic databases using keywords related to "dynamic difficulty adjustment," "adaptive game AI," "reinforcement learning in games," "player modeling in game AI," "personalized game play," and "player modeling in game AI." The sources were included based on their relevance to "difficulty adaptation in game AI," "game AI architectures," "player-centric game play difficulty adaptation," and "player-centric game play challenge adaptation."

Finally, these sources were classified into four main AI techniques used in game difficulty management: Finite State Machine, Behavior Tree, Reinforcement Learning, and Goal-Oriented Action Planning. For each of these techniques, information is extracted in terms of how each one represents game state and player state, how each one makes decisions regarding game difficulty, and what each one has to offer in terms of game engagement and flow.

The reinforcement learning-based approaches were further analyzed in detail. The benefits of the approaches were reviewed in terms of how games are formulated in Markov Decision Processes, state and action space formulation for

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 03 , March - 2026

difficulty adjustment, and reward function formulation for maintaining players within a specific performance level. The challenges of the approaches, such as quality of adaptation, computational costs, and training, were also noted. Finally, the insights from all the categories of approaches were synthesized for comparing traditional rule-based approaches with reinforcement learning-based adaptive difficulty adjustment. The synthesis of insights from all the approaches was used to highlight the benefits and challenges of the approaches, gaps in current practice, and considerations for future game development in terms of technical, design, and ethical issues related to reinforcement learning-based adaptive difficulty adjustment.

7. PROPOSED METHODOLOGY

The methodology proposed provides guidelines on how the adaptive difficulty control using reinforcement learning can be conceptualized and evaluated in the context of digital games. It is based on a series of steps starting with the definition of the problem and ending with the evaluation, without the need to develop the actual application.

In the first place, the difficulty control is conceptualized as a sequential decision-making process. The objective is to maintain the player within the desired performance level, also called the flow level, avoiding both boredom and frustration. It is important to establish the objectives in terms of the success rates, times, and errors within the desired level.

In the next step, the context in which the digital game is situated is established. It is important to establish the scenario in which the parameters related to the difficulty can be controlled. It can be related to the difficulty level of the enemies, the puzzles, the availability of the resources, among other factors. It is important to establish the state space in which the performance of the player is evaluated. It includes the performance level and the trends related to the performance, which can be established based on the player's performance.

The space of actions and difficulty parameters is then defined. The actions can vary from adjusting the health of enemies, adjusting the damage of enemies, modifying the puzzle difficulty, or modifying the resource drops. All these actions are constrained within a certain range to avoid drastic changes in the difficulty level. At the same time, a reward function is also defined. The reward function should encourage the player to stay within the target performance band. At the same time, it should discourage repeated failures or trivial successes. Smoothness can also be incorporated into the reward function to avoid abrupt changes in the difficulty level by discouraging repeated changes in the difficulty level by the agent. Depending on the state, actions, and reward functions defined, a suitable reinforcement learning technique can be chosen. The reinforcement learning techniques can vary from simple Q-learning for simple environments to deep reinforcement learning for more complex environments. Initially, the agent can be trained by simulating players or by using recorded data of players. The agent can then be integrated into the game environment. The agent can observe the state of the game and adjust the difficulty level between encounters, levels, or waves.

Finally, there is the evaluation and refinement step in the methodology. At this step, the RL-based adaptive difficulty is compared and refined against fixed difficulty modes or simple rule-based adjusters based on various metrics, including success rate, time taken, retry count, and subjective engagement level wherever possible. Observations of problems in the system, including difficulties in maintaining a balance and fairness in difficulty and transparency in decision-making, are used to refine the state representation, rewards, and granularity of actions in a way that is mindful of implementation realities in terms of computational overhead and the need for transparent and predictable systems in game development pipelines.

8. CONCLUSION

This research aims to explore adaptive difficulty adjustment in digital games in terms of reinforcement learning as a possible foundation for game personalization. The discussion in this part highlighted how game developers commonly use traditional game programming techniques like Finite State Machine, Behavior Tree, and GOAP in game development in terms of game difficulty adjustment; however, these techniques were recognized as having limitations in accommodating different player skills and play styles in modern gaming environments. These techniques were observed as having a high possibility of moving away from a player's optimal level of difficulty in a game due to their fixed and coarse-grained nature.

On the contrary, reinforcement learning has been recognized as a more adaptive and flexible game programming technique in game development in terms of game difficulty adjustment by learning from players' performance in a game and adapting game parameters in real-time. It has been recognized as having a high possibility of moving game systems closer towards a game design philosophy in which game difficulty is dynamically adjusted in accordance with a player's skill level in a game in a way that a player can feel a state of flow in a game.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 03 , March - 2026

At the same time, the study highlighted the importance of several limitations and challenges. Reinforcement learning approaches can be computationally intensive, dependent on reward functions, and liable to unforeseen behavior if not well controlled. Other factors to consider include those related to transparency, equity, and trust among players. These factors, among others, underscore the importance of careful design and control, as well as the need to communicate effectively with players. Overall, the results support the argument that the incorporation of reinforcement learning into the management of difficulty levels in games is indeed a promising path, although one to be pursued with equal vigor. ethical responsibility.

9. REFERENCES

[1] Paraschos, P. D., & Koulouriotis, D. E. (2023). Game difficulty adaptation and experience personalization: A literature review. *International Journal of Human–Computer Interaction, 39*(1), 1–22.

[2] Zhang, Y., & Goh, W. B. (2021). Personalized task difficulty adaptation based on reinforcement learning. *User Modeling and User-Adapted Interaction, 31*(4), 753–784.

[3] Mortazavi, F., Moradi, H., & Vahabie, A. H. (2024). Dynamic difficulty adjustment approaches in video games: A systematic literature review. *Multimedia Tools and Applications, 83*(35), 83227–83274.

[4] Lopes, R., & Bidarra, R. (2011). Adaptivity challenges in games and simulations: A survey. *IEEE Transactions on Computational Intelligence and AI in Games, 3*(2), 85–99.

[5] Zohaib, M. (2018). Dynamic difficulty adjustment (DDA) in computer games: A review. *Advances in Human–Computer Interaction, 2018*(1), 5681652.