

Dynamic difficulty adjustment using deep reinforcement learning: A review

Tianyi Zheng

College of Computer Science and Software Engineering, Hohai University, China

2385455009@qq.com

Abstract. With the evolution of the gaming industry, there is an increasing demand for games that offer immersive experiences beyond basic gratification. Traditional games with fixed difficulty levels often fail to cater to the diverse preferences and skills of players. To address these needs, dynamic difficulty adjustment (DDA) has emerged as a crucial element in game design. This review explores the application of Deep Reinforcement Learning (DRL) in achieving DDA, contrasting it with traditional methods. By examining various DRL algorithms and their effectiveness, as well as evaluating traditional models, this review identifies the potential of DRL in enhancing player experiences and outlines existing challenges. It also highlights future research directions, including the integration of DRL with Flow Theory and innovative evaluation methods. The goal is to provide a comprehensive overview of how DRL can advance game design and improve player satisfaction.

Keywords: Dynamic Difficulty Adjustment (DDA), Deep Reinforcement Learning (DRL), Game Design, Flow, Player Experience.

1. Introduction

With the advancement of the modern gaming industry, players increasingly seek games that offer a more immersive experience, encompassing emotional and spiritual engagement beyond mere satisfaction of basic impulses. Traditional games with static difficulty curves often fail to accommodate the diverse preferences and skill levels of players.

To meet these evolving expectations, dynamic difficulty adjustment has been facilitated using deep learning models and related algorithms. Among these, Deep Reinforcement Learning (DRL)—a subset of Deep Learning—stands out for its ability to diversify game strategies and adapt to a wide range of scenarios in innovative ways, as compared to traditional models [1]. However, DRL also faces its own set of challenges, and its effectiveness in enhancing player experience from a game design perspective necessitates further evaluation.

This review explores the integration of deep reinforcement learning with dynamic difficulty adjustment, while also considering research utilizing traditional models. It provides an examination of various algorithms and applications within this domain, evaluates their outcomes from a game design methodology perspective, identifies existing challenges, and proposes future research directions. The aim is to deliver a comprehensive overview of how deep reinforcement learning can contribute to advancing modern game design at the implementation level.

2. Background on Dynamic Difficulty Adjustment

2.1. *The importance of dynamic difficulty adjustment*

As a primary entertainment modality in contemporary life, video games offer unparalleled recreational experiences. Players are not just participants executing game rules but recipients of sensations and experiences crafted by the game. Game difficulty, an inherently subjective metric, often determines the quality of gameplay and overall experience [2]. This difficulty encompasses factors such as resource availability, understanding of game mechanics, opponent skill levels, success probabilities, and mission completion rates. Effective games often integrate tailored difficulty levels to swiftly address players' psychological needs, enhance immersion, and sustain motivation.

However, individual player preferences vary significantly, which poses a challenge for traditional games with static difficulty settings to engage a diverse player base simultaneously. To address this, the concept of Dynamic Difficulty Adjustment (DDA) has emerged [3]. DDA aims to offer players of varying skill levels a more inclusive experience by aligning gameplay with optimal flow curves [4]. This approach not only refines the game experience but also helps players quickly enter and maintain their ideal state of play [5][6].

Moreover, DDA maintains game challenge while reducing players' frustration from overly difficult challenges. This approach highlights other critical game elements such as aesthetics, narrative, music, and mechanics. Essentially, DDA represents a strategic evolution in game design, ensuring that the experience remains engaging, inclusive, and optimized for each individual player.

2.2. *Different strategy for dynamic difficulty adjustment in game design*

In contemporary game design, several approaches to adjusting game difficulty have gained prominence. These strategies include the following [7]:

(i) **Progressive Difficulty Increase upon Success:** This fundamental approach involves escalating game challenges as players progress. As characters (or players themselves) develop throughout the game, the difficulty naturally increases to sustain engagement. However, a fixed incremental difficulty can lead to monotony and may not cater to the diverse flow needs of different players.

(ii) **Pacing Player Progression:** This strategy ensures that the time required to overcome simplistic enemies does not equal that needed for more formidable ones. By adapting game difficulty to the player's skill level, this approach helps players quickly enter the optimal segment of the game, enhancing their overall experience. It significantly reduces the risk of skilled players abandoning the game during its early stages and facilitates a smoother transition into the flow state, thereby minimizing player attrition.

(iii) **Preset Difficulty Levels:** Some games allow players to select predefined difficulty tiers (e.g., Easy, Normal, Hard, Hell) at the beginning, which are then applied consistently throughout the gameplay. This method enables players with prior experience to autonomously choose a difficulty level that best suits their preferences. Essentially, this strategy implements Dynamic Difficulty Adjustment (DDA) from the player's perspective.

(iv) **Incorporating Respite and Reinforcements:** To alleviate anxiety from challenging scenarios, some games include mid-game respites or reinforcements that help players recover and regain their composure. This approach aims to guide players back toward an optimal flow state by providing moments of relief and restoration.

2.3. *Traditional DDA approaches*

Over the past two decades, numerous researchers have made significant strides in the implementation of Dynamic Difficulty Adjustment (DDA) [8]. Various innovative models and algorithms have been employed to analyze players' in-game performance, classify their skill levels, and adjust game difficulty accordingly.

Khajah et al. [9] applied Bayesian Optimization to adjust player inputs in two action-based games, *Flappy Bird* and *Spring Ninja*. This approach facilitated real-time game control adjustments and

achieved notable results in DDA. However, extending this methodology to more complex scenarios remain an area ripe for further investigation.

More recently, Romero-Mendez et al. [10] employed a Feedforward Neural Network (FNN) model within deep learning to classify and predict player skill levels in a Space Invaders-like arcade game. Based on these predictions, a DDA strategy tailored to predefined difficulty tiers was implemented, yielding promising outcomes.

Sutanto and Suharjito [11] enhanced Artificial Neural Networks (ANNs) by incorporating the Adaptive Neuro-Fuzzy Inference System (ANFIS). This approach analyzed parameters such as enemy kills and hero damage in a Shoot 'Em Up (STG) game, categorizing players into five distinct difficulty levels.

It is evident that most traditional DDA approaches focus on player-level adjustments, using predefined rules to broadly categorize player abilities and adjust difficulty. These methods typically ensure high prediction accuracy and effective model performance. However, such approaches may struggle with complex situations or capturing real-time player states across multiple dimensions due to their reliance on heuristic rules. Additionally, these methods often require the prior construction of mathematical models tailored to specific games, necessitating manual updates if game elements undergo significant changes.

3. Approaches of Applying Dynamic Difficulty Adjustment in Game Design Area

3.1. Related works using deep reinforcement learning models

To explore alternative approaches to Dynamic Difficulty Adjustment (DDA), some researchers have turned to Deep Reinforcement Learning (DRL) models. Reinforcement Learning (RL) has already demonstrated significant success in managing complex game strategies, and DRL models extend these capabilities by dynamically adjusting various difficulty parameters in real time. This allows for adaptive gameplay difficulty tailored to individual players, with DRL models consistently achieving impressive results.

Several researchers have employed DRL models to create diverse difficulty levels within games and maintain player flow. For instance, Huber applied DRL in VR sports games to dynamically generate game levels based on player performance, thus enriching gameplay and enhancing player motivation to continue playing [12].

Additionally, the intensity of AI opponents in games significantly affects the operational difficulty and strategic depth of gameplay. Some researchers have focused on real-time adjustment of AI opponent intensities. Wang et al. [13] proposed using DRL models to develop a smart agent that adjusts game difficulty according to player proficiency. This agent, which integrated DRL with action selection mechanisms, was effectively used in real-time fighting games to enable adaptive AI opponent strategies. Furthermore, Wender and Watson [14] addressed city placement tasks in the SLG game *Civilization VI*, creating more challenging AI opponents. Sutton et al. [15] and Climent et al. [16] applied the SARSA algorithm in *Space Invaders* to achieve DDA, designing a DRL agent that yielded commendable results in turn-based single-player games.

Despite these promising outcomes, DRL models are known to be sensitive to hyperparameters. Variations in hyperparameters can lead to substantial performance differences even within the same algorithmic framework. Moreover, the reinforcement learning process carries the risk of deviating from intended goals. To address these challenges, DRL with Evolutionary Algorithms (EA) [17][18] are integrated and introduced. The REAs algorithm, which is specifically designed for single-objective optimization [19], has demonstrated notable advantages over traditional DRL models in learning capability, model convergence speed, and execution efficiency. It effectively mitigates issues such as sparse reward problems, overestimation of Q-values by Q-Learning networks [20][21], and insufficient exploration. Additionally, RMOEA algorithm excels in addressing multi-objective optimization problems.

3.2. Model-related challenges and evaluation challenges

Implementing DDA through DRL models presents several significant challenges. Firstly, to address the risk of overfitting, it is crucial to involve players with varying skill levels during the development phase. This practice enables the model to adapt more effectively to a range of player abilities and environments, thereby enhancing its robustness and generalizability.

Secondly, the requirements for DDA can vary considerably across different game genres. For example, in puzzle-based games like *Angry Birds*, DDA might involve generating levels with varying difficulty based on the player's proficiency. Conversely, in first-person shooters (FPS), DDA could focus on adjusting variables such as the number of enemies or the damage inflicted. In Multiplayer Online Battle Arena (MOBA) games, DDA might modulate player skill levels during matchmaking or ranking processes [22].

From a game design perspective, the implementation of DDA through DRL models can sometimes be perceived as monotonous or disengaging. Simply providing players with an abundance of resources may lead to decreased motivation, while excessive reliance on probabilistic adjustments might undermine the sense of challenge.

Moreover, optimizing game experience through DRL models necessitates evaluation from the player's perspective, but the metacognitive Heisenberg effect poses a paradox in assessing DDA [23]. Specifically, players need to evaluate whether DDA improves their gameplay experience, yet their focus on DDA mechanisms during gameplay can disrupt their flow state, as individuals often struggle to analyze their experiences while they are occurring.

3.3. Possible solutions

Ultimately, not all games are suitable for DDA, and no single model or algorithm can universally adapt to every game type. Therefore, DDA should be applied judiciously and with a clear purpose tailored to the specific context of the game. When choosing appropriate models, game designers must possess a deep understanding of their game's unique characteristics to ensure that each difficulty adjustment—whether static or dynamic—serves a meaningful purpose and enhances the overall player experience.

To address evaluation challenge, Andrade et al. [24] has identified three fundamental requirements for implementing DDA from a game design perspective. Additionally, Jesse Schell advocates for a methodical approach where players first experience the game and subsequently analyze it separately [7]. These strategies provide valuable guidance in assessing the true impact of DDA. For example, a two-playthrough approach, where players engage with the game twice, can be employed. Moreover, incorporating subtle questions in post-game surveys, such as "Did you perceive your actual gameplay time as longer or shorter than anticipated?" can offer indirect insights into the player's flow state.

4. Future Direction

The journey toward realizing DDA through DRL models requires further exploration in several key areas. Firstly, a theoretical foundation is needed to fully leverage the advantages of DRL models and apply them effectively across a diverse range of game types. This includes developing a comprehensive understanding of how DRL can be tailored to various gaming contexts.

Secondly, although Flow Theory may present constraints on DDA designs [25], the intersection of DDA and Flow Theory remains a crucial research direction. It is essential to clarify whether DRL models can elucidate the process of enhancing player experience by effectively analyzing and optimizing players' flow states.

Additionally, innovative approaches to DDA involving DRL models warrant consideration. For instance, integrating electroencephalography (EEG) [26], physiological measures, and subjective ratings [27] could offer novel ways to gather valuable data, thereby advancing research in this field. Establishing normative and unambiguous evaluation criteria is also necessary to accurately assess whether a model enhances the player experience.

5. Conclusion

DRL presents a promising avenue for advancing DDA in video games. Compared to traditional methods, DRL offers greater flexibility and adaptability, potentially leading to more engaging and personalized player experiences. However, the implementation of DRL-based DDA is not without its challenges, particularly concerning game design and model performance. Effective evaluation of DRL models requires a player-centric approach to ensure that the game experience is optimized. Future research should focus on harnessing the full potential of DRL across various game types, exploring its relationship with flow theory, and establishing clear, normative criteria for evaluating DDA effectiveness.

References

- [1] Mohammed, M., Khan, M. B., & Bashier, E. B. M. (2016). *Machine learning: algorithms and applications*. CRC Press.
- [2] Koster, R. (2014). *A Theory of Fun for Game Design*. O'Reilly Media.
- [3] Hunnicke, R. (2005). The case for dynamic difficulty adjustment in games. *Proceedings of the ACM SIGCHI International Conference on Advances in Computer Entertainment Technology (ACE '05)*, June 2005, Valencia, Spain, 429–433.
- [4] Csikszentmihalyi, M. (2009). *Flow: The Psychology of Optimal Experience*. Harper Row.
- [5] Chen, J. (2007). Flow in Games (and Everything Else). *Communications of the ACM*, 50, 31–34.
- [6] Koster, R. (2014). *A Theory of Fun for Game Design*. O'Reilly Media.
- [7] Schell, J. (2008). *The Art of Game Design: A Book of Lenses*. CRC Press.
- [8] Zohaib, M. (2018). Dynamic Difficulty Adjustment (DDA) in Computer Games: A Review. *Advances in Human-Computer Interaction*, 1–12.
- [9] Khajah, M. M., Roads, B. D., Lindsey, R. V., Liu, Y.-E., & Mozer, M. C. (2016). Designing engaging games using Bayesian optimization. *Proceedings of the 34th Annual Conference on Human Factors in Computing Systems (CHI 2016)*, May 2016, San Jose, CA, USA, 5571–5582.
- [10] Romero-Mendez, E. A., Santana-Mancilla, P. C., Garcia-Ruiz, M., Montesinos-López, O. A., & Anido-Rifón, L. E. (2023). The Use of Deep Learning to Improve Player Engagement in a Video Game through a Dynamic Difficulty Adjustment Based on Skills Classification. *Applied Sciences*, 13, 8249.
- [11] Sutanto, K., & Suharjito, S. (2014). Dynamic difficulty adjustment in games based on type of player with ANFIS method. *Journal of Theoretical and Applied Information Technology*, 10, 254–260.
- [12] Huber, T., Mertes, S., Rangelova, S., Flutura, S., & André, E. (2021). Dynamic Difficulty Adjustment in Virtual Reality Exergames through Experience-driven Procedural Content Generation. *2021 IEEE Symposium Series on Computational Intelligence (SSCI)*, Orlando, FL, USA, 1–8. doi: 10.1109/SSCI50451.2021.9660086
- [13] Wang, H., Gao, Y., & Chen, X. (2010). DRL-dot: A reinforcement learning NPC team for playing domination games. *IEEE Transactions on Computational Intelligence and AI in Games*, 2(1), 17–26.
- [14] Wender, S., & Watson, I. (2008). Using reinforcement learning for city site selection in the turn-based strategy game *Civilization IV*. *2008 IEEE Symposium on Computational Intelligence and Games (CIG)*, 372–377.
- [15] Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction*. MIT Press.
- [16] Climent, L., Longhi, A., Arbelaez, A., & Mancini, M. (2024). A framework for designing Reinforcement Learning agents with Dynamic Difficulty Adjustment in single-player action video games. *Entertainment Computing*, 50, 100686.
- [17] Kaidan, M., Harada, T., Chu, C. Y., & Thawonmas, R. (2016). Procedural generation of Angry Birds levels with adjustable difficulty. *2016 IEEE Congress on Evolutionary Computation (CEC)*, 1311–1316.

- [18] Ferreira, L., & Toledo, C. (2014). A search-based approach for generating Angry Birds levels. *Computational Intelligence & Games*, IEEE.
- [19] Cui, G., Shen, R., Chen, Y., Zou, J., Yang, S., Fan, C., & Zheng, J. (2020). Reinforced Evolutionary Algorithms for Game Difficulty Control. In *Proceedings of the 2020 3rd International Conference on Algorithms, Computing and Artificial Intelligence* (pp. 1–7).
- [20] Bai, H., Cheng, R., & Jin, Y. (2023). Evolutionary reinforcement learning: A survey. *Intelligent Computing*, 2, 0025.
- [21] Mnih, V., Kavukcuoglu, K., Silver, D., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529.
- [22] Silva, M. P., Silva, V. do N., & Chaimowicz, L. (2017). Dynamic difficulty adjustment on MOBA games. *Entertainment Computing*, 18, 103–123.
- [23] Andrade, G., Ramalho, G., Santana, H., & Corruble, V. (2005). Extending reinforcement learning to provide dynamic game balancing. *Proceedings of the Workshop on Reasoning, Representation, and Learning in Computer Games, 19th International Joint Conference on Artificial Intelligence (IJCAI)*, Edinburgh, United Kingdom, 7, 12.
- [24] Koriat, A., & Bjork, R. A. (2006). Metacognition and the monitoring of decoding processes.
- [25] Guo, Z., Thawonmas, R., & Ren, X. (2024). Rethinking dynamic difficulty adjustment for video game design. *Entertainment Computing*, 50, 100663.
- [26] Fisher, N., & Kulshreshth, A. K. (2024). Exploring Dynamic Difficulty Adjustment Methods for Video Games. *Virtual Worlds*, 3(2), 230–255.
- [27] Ozkul, F., Palaska, Y., Masazade, E., & Erol-Barkana, D. (2019). Exploring dynamic difficulty adjustment mechanism for rehabilitation tasks using physiological measures and subjective ratings. *IET Signal Processing*, 13, 378–386.