applied sciences

MDPI

Article

Dynamic Difficulty Adjustment with Machine Learning for Air Hockey

Mikhail Zgonnikov [1] and Maxim Mozgovoy [2,*]

[1] School of Engineering and Computer Science (EPITA), 16 Rue Jean-Marie Leclair, 69009 Lyon, France; zgonnikovmichka@gmail.com
[2] School of Computer Science and Engineering, The University of Aizu, Aizuwakamatsu 965-8580, Fukushima Prefecture, Japan
* Correspondence: mozgovoy@u-aizu.ac.jp

Abstract

This work presents a method for implementing dynamic difficulty adjustment in the arcade game of Air Hockey using reinforcement learning. The resulting AI-controlled opponent is capable of adapting its skill level to the player's performance to maintain engagement and provide a balanced gameplay experience. The approach relies on generating several AI agents through progressively longer training durations, resulting in distinct and smoothly transitioning difficulty levels that can be switched dynamically. We discuss how this scheme can be extended with manually selected parameters that influence physical aspects of the agent's behavior—such as movement speed, reaction latency, and control precision—to complement the variations in decision-making quality. The proposed method is applicable to a wide range of video games, and experimental results demonstrate its effectiveness in producing adaptive and varied opponent behavior.

Keywords: machine learning; reinforcement learning; dynamic difficulty adjustment; flow channel; air hockey

1. Introduction

One of the central challenges faced by video game designers is maintaining player engagement throughout the gaming experience. Engagement is often conceptualized through the notion of a Flow Channel (FC), a psychological state in which the player remains situated between anxiety and boredom, thereby sustaining motivation and enjoyment [1]. This balance is delicate: anxiety typically arises when the game becomes too difficult relative to the player's abilities, whereas boredom emerges when the challenge is insufficient. Both extremes can lead to disengagement, making it essential for games to adapt to the diverse skill levels of their audience.

A widely explored solution to this problem is the use of Dynamic Difficulty Adjustment (DDA) techniques [2]. DDA systems aim to continuously align the game's difficulty with the player's performance, ensuring that the experience remains stimulating without becoming overwhelming. Although many factors can influence perceived difficulty, the behavior and skill level of AI-controlled opponents often play a decisive role. Adjusting these opponents, whether by modifying their strategic reasoning, reaction time, or precision, has proven to be an effective method for implementing DDA in a broad range of game genres [3]. As a result, DDA has become increasingly prevalent in modern game design [4].

Academic Editor: Firstname Lastname

Received:
Revised:
Accepted:
Published:

In parallel, the growing integration of machine learning techniques into video game AI has opened new possibilities for creating adaptive and believable opponents [5]. Building on this trend, the objective of the present work is to investigate how machine learning agents, combined with DDA, can be used to construct a flexible and responsive opponent for a simple arcade-style Air Hockey game. Our approach relies on training several AI agents for different durations to produce opponents with distinct skill levels. These agents can be interchanged dynamically, either at the start of a match or in response to in-game events, such as goals scored by the player or the AI. Because the agents share similar training environments and architectures, adjacent difficulty levels exhibit consistent behavioral patterns, ensuring smooth transitions between them.

Through a series of controlled matches, we demonstrate that this method successfully generates opponents with clearly differentiated abilities, suitable for use in a DDA framework. Furthermore, we argue that the same methodology can be generalized to other game types, since it does not rely on characteristics unique to Air Hockey. Instead, it provides a modular and scalable strategy for integrating machine learning based difficulty adjustment into interactive digital environments.

2. The Game of Air Hockey

A video game adaptation of Air Hockey draws its inspiration from the real-world tabletop sport of the same name. The physical game is played by two opponents on a raised, low-friction table, typically measuring 8 feet in length and 4 feet in width. Each player uses a round mallet to strike a lightweight plastic puck with the objective of sending it through the opponent's goalposts. The characteristic low friction of the playing surface is achieved through a continuous airflow produced by fans integrated into the table, which allows the puck to hover slightly above the surface and glide with minimal resistance.

Although Air Hockey is often perceived as a simple and fast-paced dexterity game, effective play requires a combination of precision, anticipation, and tactical decision-making. Even novice players quickly encounter strategic concepts that may initially seem counterintuitive. One well-known example is the "triangular defense", a technique that instructs players to position their mallet away from the goal line rather than directly in front of it, thereby improving defensive coverage and reaction time [6]. As players develop their skills, they learn to balance offensive opportunities with defensive stability, making the game both accessible and strategically engaging.

Air Hockey as a video game shares certain similarities with classic paddle-and-ball games like Pong, but its more advanced gameplay contributes to its enduring popularity, especially on mobile platforms that allow realistic touch-based controls. Given that most Air Hockey apps are similar in terms of gameplay, graphics and AI are the factors that make a specific implementation stand out. Arguably, in single-player mode, AI is the most important factor affecting the player's Flow Channel.

3. Related Works

DDA has long been recognized as an effective strategy for enhancing player enjoyment and sustaining engagement in digital games [7]. Increased player satisfaction has been reported across a wide range of genres, including classic puzzle games such as Tetris [8], first-person shooters [9,10], various action-oriented titles [11], and even cognitive training environments such as visual working memory games [12]. As noted by Zohaib [2], a fundamental component shared by most DDA approaches is the ability to estimate the level of challenge experienced by the player at any given moment. This estimation serves as the basis for adapting game parameters to maintain player engagement while avoiding frustration or boredom. Existing research proposes a wide range of methods for this

https://doi.org/10.3390/app1010000

purpose, varying significantly in complexity. At the simplest level, challenge estimation can rely on in-game performance metrics such as score evolution, success and failure rates, or reaction times [13]. These score-based heuristics are particularly well-suited for games with clearly defined objectives and limited state spaces. Given the relatively simple and highly structured gameplay of Air Hockey, this category of methods is adopted in this work, as it enables efficient and transparent difficulty adaptation without introducing significant noise or latency.

More complex approaches attempt to infer player experience through indirect measurements. For instance, Mi and Gao [14] treat DDA as a method of churn prevention, and thus attempt to track player engagement rather than straightforward win/loss ratio. Rani et al. [15] investigate the use of physiological signals, such as heart rate and galvanic skin response, to dynamically regulate task difficulty based on player emotional state. While such techniques can provide a finer-grained estimation of stress or engagement, they require specialized hardware and are therefore less practical for standard gaming environments.

Recent works have also explored nonintrusive perceptual methods. Blom et al. [16] present a system capable of continuously adapting game difficulty during live gameplay by analyzing facial expressions. This approach allows difficulty adjustment without explicit player input or gameplay interruption, but relies on computer vision pipelines that may be sensitive to environmental conditions and computational constraints.

Overall, these studies highlight a trade-off between the accuracy of player state estimation and the practical feasibility of deployment. In the context of this project, score-based evaluation was selected as a robust and lightweight solution that aligns with the fast-paced nature of Air Hockey and integrates naturally with the Reinforcement Learning (RL)-based opponent design.

Once the system identifies the need for adjustment, the implementation of difficulty changes typically relies on highly game-specific mechanisms. Because game mechanics vary significantly across genres, it is challenging to define universally applicable methods. Hossan et al. [17] discuss several notable approaches, including dynamic scripting, fast Bayesian content adaptation, and RL. Still, most studies focus on tailoring the difficulty modulation process to the unique characteristics of the game under consideration.

The task of developing AI opponents for Air Hockey has also been explored in prior work. Some physical robotics-based projects, such as [18], investigated mathematically grounded strategies for controlling an Air Hockey robots. Other studies leveraged machine learning to create more flexible and adaptive agents, as demonstrated in [19], although these efforts did not incorporate DDA.

The application of DDA specifically to Air Hockey was examined by Delgado-Mata and Ibáñez [20]. Their approach is distinctive in that it focuses on modifying the game environment rather than altering the skill of an AI-controlled opponent. This method is particularly suitable for player-versus-player scenarios, where both participants are human. The authors employ a score-based heuristic to adjust the width of each player's goal line and to modify the friction of the playing surface, thereby influencing the likelihood of scoring.

More broadly, RL and artificial intelligence have become increasingly prominent in the design of video game agents, in particular, in classic and mechanically simple game environments. These technologies offer significant potential for enhancing DDA systems and maintaining player FC. Several studies have demonstrated promising results in this direction [21–24], highlighting the capacity of AI-driven methods to create adaptive and personalized gameplay experiences.

https://doi.org/10.3390/app1010000

4. DDA in Air Hockey

4.1. 3D-AirHockey

The computer simulation of Air Hockey used as the testbed for this project is the open-source game 3D-AirHockey, developed in Unity by Andrei Lapusteanu (https://github.com/Andrei-Lapusteanu/3D-AirHockey (accessed on 17 March 2026)). The game provides a local two-player experience designed for casual, "couch-style" play and is presented through a top-down camera view (see Figure 1). Each player controls a mallet using four directional buttons, while mouse input is intentionally disabled. A mallet gradually accelerates when moved continuously in the same direction, which introduces a degree of challenge for human players who must adapt to the momentum-based controls. This design choice also simplifies the implementation of an AI controller, as simulating discrete button presses is considerably more straightforward than replicating continuous mouse movements.

For the purposes of this project, we created a fork of Lapusteanu's original game (https://github.com/Net-Runer/AirHockeyProject (accessed on 17 March 2026)), integrating a custom AI controller while preserving all other gameplay elements, physics, and visual components. This ensures that the behavior of the AI agents can be evaluated within an environment that remains faithful to the original design, allowing for consistent comparisons between human and machine-controlled gameplay.

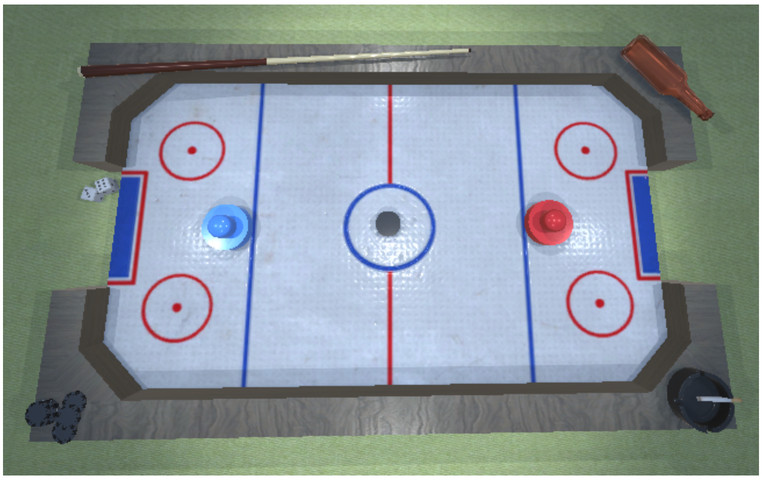

Figure 1. A screenshot of 3D-AirHockey.

4.2. DDA via Reinforcement Learning

RL is a subfield of machine learning that studies how an autonomous agent can learn to make sequential decisions by interacting with an environment in order to maximize its cumulative reward. At each time step, the agent observes the current state of the environment, selects an action according to a policy, and receives a scalar reward that evaluates the quality of that action. The environment then transitions to a new state according to its underlying dynamics. This interaction loop is commonly formalized using the framework of Markov Decision Processes, which define the set of states, available actions, transition probabilities, and reward functions. Through repeated trial and error, and by balancing exploration of unfamiliar actions with exploitation of known beneficial ones, the agent gradually improves its policy.

In 3D-AirHockey, our goal is to apply RL and DDA to obtain an AI-controlled opponent able to facilitate player FC. Since the original project only supported local two-player matches, we had to extend it with an AI controller functionality. Our RL-based system generates a collection of dynamically switchable AI opponents, representing distinct skill levels.

To introduce RL capabilities into the game, we used Unity Machine Learning Agents Toolkit (ML-Agents) (https://github.com/Unity-Technologies/ml-agents (accessed on 17 March 2026)), which provides convenient access to a wide range of machine learning algorithms, similar to those used in [5]. For training, we selected the proximal policy optimization (PPO) algorithm and employed a multi-layer perceptron network composed of two hidden layers with 256 neurons each. Both players were controlled by AI agents during training, enabling RL through self-play. The AI controller emulates human input by producing two output values corresponding to horizontal and vertical directional keypresses (see Figure 2).

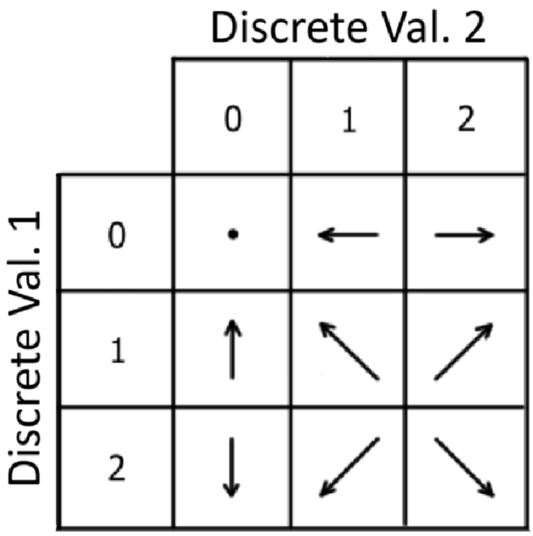

Figure 2. Possible output values of AI controller.

The agent receives the following sensory information as input:

- Positions of the AI bot and its opponent (3D coordinates).
- Velocities of the AI bot and its opponent (3D directional vectors).
- Position of the puck (3D coordinates).
- Velocity of the puck (3D directional vector).

Although it is possible to design a reward function based solely on goals scored, faster training time and more meaningful intermediate strategies can be obtained with additional rewards, based on the concepts of *defense zone* and *playing zone* (see Figure 3). Defense zone corresponds to the area immediately surrounding the player's goal. Playing zone denotes the whole player's half of the game field. The idea is to encourage the AI player to remain close to its defense zone when the puck is located outside of its playing zone.

Training was carried out in a parallelized setup, using 100 subprocesses running simultaneously. To maximize training efficiency and reduce computational overhead, all nonessential visual components of the environment, including lighting effects and decorative elements, were disabled. This configuration allowed the learning process to focus exclusively on gameplay mechanics and agent decision making rather than rendering.

The reward function was designed to guide the agent toward effective and realistic Air Hockey strategies. The full set of rewards and penalties is presented in Table 1. Rewards marked with $+D$ are computed as a function of the distance between relevant entities, such as the player and the puck, and are defined as follows:

$$D = 0.1 \times Clamp_{[0,1]}\left(1 - \frac{distance}{5}\right) \tag{1}$$

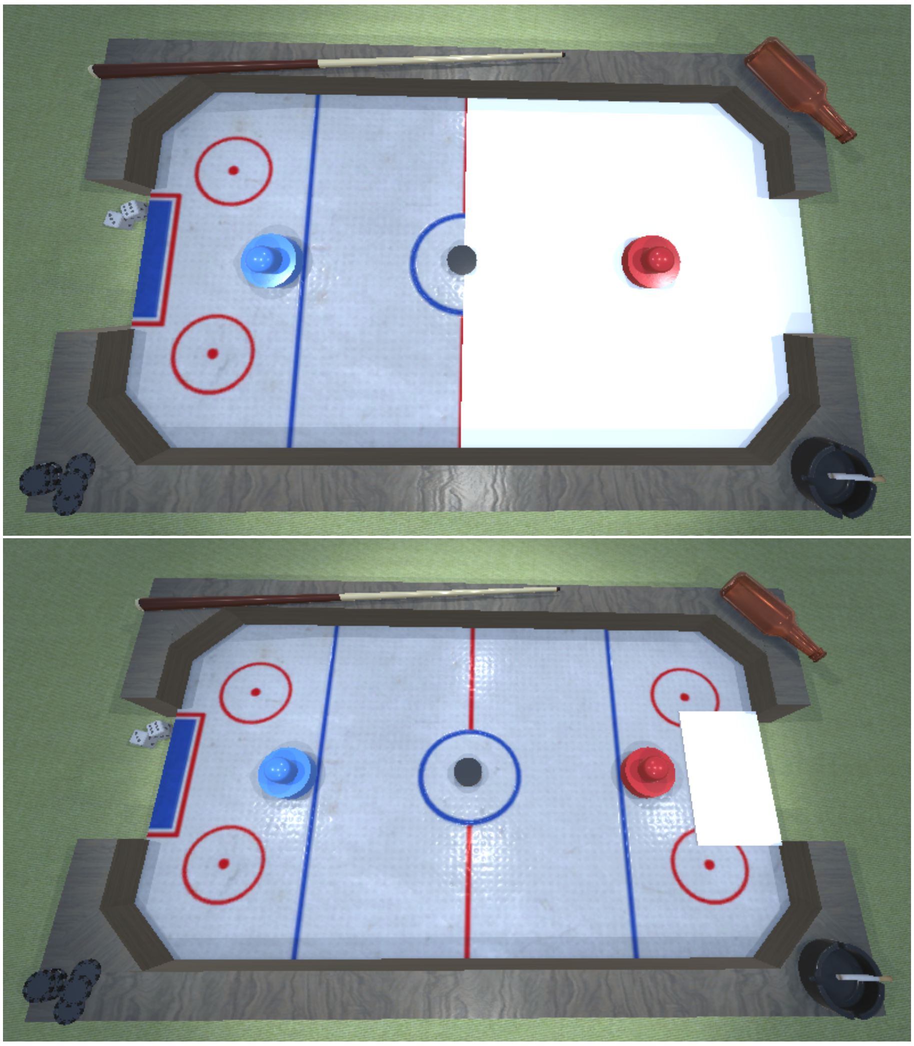

Figure 3. AI-controlled areas: playing zone (**above**) and defense zone (**below**).

This formula provides a smooth incentive that increases as the distance decreases, while remaining bounded. The clamping function ensures numerical stability and prevents extreme reward values:

$$Clamp_{[0,1]}(x) = \begin{cases} 0, & if\ x < 0, \\ x, & if\ 0 \leq x \leq 1, \\ 1, & if\ x > 1. \end{cases} \tag{2}$$

A specific negative reward was introduced when the agent touches the puck while its velocity is close to zero. This penalty addresses the issue known as *reward farming*, where an agent exploits unintended loopholes in the reward structure to accumulate rewards without achieving the desired objective. In our case, such behavior occurs when the puck is trapped between the mallet and a wall, allowing the agent to repeatedly trigger the puck contact reward without performing a meaningful action. The penalty discourages this exploit while preserving the original intention of the contact reward, which is to motivate active puck engagement rather than passive positioning or avoidance.

Overall, the reward structure balances offensive incentives, such as scoring goals and shooting toward the opponent's side, with defensive behaviors, including positioning between the puck and the goal and controlling puck movement in the defensive zone. Minor penalties were also applied to discourage unrealistic or ineffective behaviors, such as excessive wall contact, idling, or allowing the puck to pass behind the player.

https://doi.org/10.3390/app1010000

Table 1. Rewards and Penalties.

Event	Value
Scoring a goal	+4
Conceding a goal	−8
Touching the puck	+1
Staying close to the puck	$+D$
Touching a wall	−0.1
Shooting the puck toward the opponent's goal	+0.1
Shooting the puck away from the opponent's goal	−0.1
Allowing the puck to pass behind the player	−0.02
Staying between the puck and the goal	+0.05
Aligning the player, the puck, and the opponent's goal	+0.05
Staying close to the defense zone when the puck is out of playing zone	$+D$
Keeping the puck velocity low	−0.005
Touching the puck when its speed is near zero	−0.05
Idling	−0.001

The AI controller reaches a level of performance corresponding to a reasonably skillful human player after approximately 5,000,000 training steps (about half of an hour of training on a conventional laptop). For the purposes of DDA, we saved the snapshots of the model after every 330,000 steps of training, thus obtaining 15 different models, presumably exhibiting different levels of skill. It turned out that the least trained model was not able to perform even the most basic game actions, so our resulting list consists of 14 models (difficulty levels 1 to 14).

The mechanism of DDA used in our project operates according to the following scheme. The game starts at the "medium" difficulty level 7. Once a goal is scored, we check if either of the players leads by at least two points. If the human player leads, we increase the current difficulty level. If the bot leads, we decrease it.

Note that the aim of this simple method is to prove that it is possible to obtain the desired skill level by switching between the models trained on different amount of data. We do not attempt to maximize player engagement (as Mi and Gao [14]), and we do not propose that locating opponents having the closest matching skill level is desirable in practice.

5. Results

While experimenting with the obtained system, we aimed to show that (1) AI indeed gains skills as a result of reinforcement learning and (2) the proposed DDA system is able to produce an AI-controlled opponent roughly equal in skill to the current human player.

To demonstrate the effect of reinforcement learning, we conducted a series of matches between AI-controlled opponents of different skill levels. Only even-numbered levels were used to reduce the resulting printout (see Table 2). Reported values are scores obtained after 10 minutes of play. Clearly, higher-skilled AIs were always able to defeat opponents having less training data. However, the difference in goals scored does not reflect the difference in skills in these short sessions.

To demonstrate the effect of dynamic difficulty adjustment in action, we asked five human subjects with diverse arcade game skills to play for 6 minutes against an AI opponent. Every game was preceded by a 6-minute training session, giving participants the chance to familiarize themselves with the game. The final scores (see Table 3) show roughly the same skill level for both opponents in every match. Interestingly, the current DDA scheme gives a slight advantage to the AI-controlled player, who won most games, albeit with a minimal margin. The table also shows that the human players possess different skill levels.

For example, H_2 matches an AI of level 3 or 4, while H_4 mostly competes against an AI of level 7.

Table 2. Results of matches between the models L_2 to L_{14}.

	L_2	L_4	L_6	L_8	L_{10}	L_{12}	L_{14}
L_2		14:39	27:41	14:43	13:56	16:54	17:37
L_4			28:36	28:42	22:43	24:52	31:47
L_6				30:42	27:37	20:43	28:44
L_8					27:35	31:38	18:36
L_{10}						32:43	21:38
L_{12}							27:39
L_{14}							

Table 3. Results of matches between human subjects H_1 to H_5 and AI-controlled opponents.

Subject	Difficulty Level at						Final Score
	01:00	02:00	03:00	04:00	05:00	06:00	
H_1	4	5	6	6	3	3	28:31
H_2	8	5	4	3	4	3	16:18
H_3	5	3	3	4	7	5	25:26
H_4	1	7	7	5	7	8	25:28
H_5	7	3	7	13	14	14	42:31

6. Discussion

The primary objective of this study was to design an adaptive AI-controlled opponent for a video game version of Air Hockey. The experimental results indicate that this objective has been successfully achieved. The 14 pre-trained AI exhibit clearly distinguishable and progressively stronger skill levels. The system is capable of switching between these models in response to changes in the game score, allowing for smooth transitions that help maintain an appropriate level of challenge throughout a match.

The proposed approach has potential applicability beyond the specific context of Air Hockey. The underlying idea of generating multiple AI opponents by varying training duration is general and can be transferred to other game genres. Because the resulting models represent incremental refinements of the same fundamental behavior, transitions between difficulty levels are likely to appear natural rather than abrupt, which is essential for preserving player immersion.

Nevertheless, the method also has limitations. The differences between the AI models arise primarily from variations in decision-making quality rather than from changes in physical attributes. In many games, characters possess distinct profiles defined by attributes such as speed, strength, stamina, or precision. In Air Hockey, such attributes may include top movement speed, minimal reaction time, and accuracy of shots. Arguably, AI-controlled characters should differ in both decision making and physical abilities. However, physical abilities are inherently game-specific, and their automated adjustment is not easy to generalize. The separation between "RL-based" and "attribute-based" behavior is elaborated in the work by Dziedzic [25], where games are classified across "pace", "topography", and "mutability" dimensions, suggesting a particular approach to DDA for the given combination.

7. Conclusions

We have demonstrated that a combination of manual parameter adjustment and machine learning techniques can be effectively employed to construct a DDA system for the game of Air Hockey. Although the resulting implementation contains several game-specific components, the overall methodology remains broadly applicable. The idea of generating multiple AI opponents through variations in training time, and possibly refining their behavior through targeted parameter modifications, can be extended to a wide range of video games, including more complex genres such as Multiplayer Online Battle Arena (MOBA) titles [10]. This type of adaptive functionality is increasingly relevant in the modern game industry, which has evolved into a major entertainment sector [26] with a highly diverse player base and rising expectations for personalization and high-quality interactive experiences [27].

Several options for future research emerge from this work. One possible direction is to examine the contribution of physical player parameters into game difficulty, and introduce parameter tuning into the DDA system. Another interesting extension would be to investigate alternative triggers for difficulty adaptation beyond goal events, e.g., connected to player engagement [14]. Such triggers may be particularly valuable in more complex games, where player performance and enjoyment can be assessed through a richer set of behavioral indicators.

Author Contributions: Conceptualization, Mikhail Zgonnikov and Maxim Mozgovoy; Methodology, Mikhail Zgonnikov and Maxim Mozgovoy; Software, Mikhail Zgonnikov; Validation, Mikhail Zgonnikov; Formal analysis, Maxim Mozgovoy; Investigation, Mikhail Zgonnikov and Maxim Mozgovoy; Writing—original draft, Mikhail Zgonnikov; Writing—review & editing, Mikhail Zgonnikov and Maxim Mozgovoy; Supervision, Maxim Mozgovoy. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original data presented in the study are openly available in repository AirHockeyProject at https://github.com/Net-Runer/AirHockeyProject.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Cowley, B.; Charles, D.; Black, M.; Hickey, R. Toward an Understanding of Flow in Video Games. *Comput. Entertain.* **2008**, *6*, 1–27.
2. Zohaib, M. Dynamic Difficulty Adjustment (DDA) in Computer Games: A Review. *Adv. Hum.-Comput. Interact.* **2018**, *2018*, 1–12.
3. Silva, M.P.; Silva, V.d.N.; Chaimowicz, L. Dynamic Difficulty Adjustment through an Adaptive AI. In *Proceedings of the 2015 14th Brazilian Symposium on Computer Games and Digital Entertainment (SBGames)*; IEEE: New York, NY, USA, 2015; pp. 173–182.
4. Mortazavi, F.; Moradi, H.; Vahabie, A.H. Dynamic difficulty adjustment approaches in video games: A systematic literature review. *Multimed. Tools Appl.* **2024**, *83*, 83227–83274.
5. Urmanov, M.; Alimanova, M.; Nurkey, A. Training unity machine learning agents using reinforcement learning method. In *Proceedings of the 2019 15th International Conference on Electronics, Computer and Computation (ICECCO)*; IEEE: New York, NY, USA, 2019; pp. 1–4.
6. Billiards, H. Professional Air Hockey: Tips, Techniques & Strategy. 2025. Available online: https://hallmarkbilliards.com/mastering-air-hockey-techniques-and-strategies-from-the-world-championships (accessed on 17 March 2026).
7. Moon, H.S.; Seo, J. Dynamic difficulty adjustment via fast user adaptation. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*; 2020; pp. 13–15.
8. Lora, D.; Sánchez-Ruiz, A.A.; González-Calero, P.A.; Gómez-Martín, M.A. Dynamic Difficulty Adjustment in Tetris. In Proceedings of the FLAIRS, Key Largo, FL, USA, 16–18 May 2016; pp. 335–339.
9. Knorr, J.; Vaz De Carvalho, C. Using Dynamic Difficulty Adjustment to Improve the Experience and Train FPS Gamers. In Proceedings of the Ninth International Conference on Technological Ecosystems for Enhancing Multiculturality (TEEM'21), Barcelona, Spain, 26–29 October 2021; pp. 195–200.

10. Ye, D.; Liu, Z.; Sun, M.; Shi, B.; Zhao, P.; Wu, H.; Yu, H.; Yang, S.; Wu, X.; Guo, Q.; et al. Mastering complex control in moba games with deep reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*; 2020; Volume 34, pp. 6672–6679.
11. Climent, L.; Longhi, A.; Arbelaez, A.; Mancini, M. A framework for designing Reinforcement Learning agents with Dynamic Difficulty Adjustment in single-player action video games. *Entertain. Comput.* **2024**, *50*, 100686.
12. Rahimi, M.; Moradi, H.; Vahabie, A.h.; Kebriaei, H. Continuous reinforcement learning-based dynamic difficulty adjustment in a visual working memory game. *arXiv* **2023**, arXiv:2308.12726.
13. Fisher, N.; Kulshreshth, A.K. Exploring dynamic difficulty adjustment methods for video games. *Proc. Virtual Worlds* **2024**, *3*, 230–255.
14. Mi, Q.; Gao, T. Engagement-Oriented Dynamic Difficulty Adjustment. *Appl. Sci.* **2025**, *15*, 5610. https://doi.org/10.3390/app15105610.
15. Rani, P.; Sarkar, N.; Liu, C. Maintaining optimal challenge in computer games through real-time physiological feedback. In *Proceedings of the 11th International Conference on Human Computer Interaction*; 2005; Volume 58, pp. 22–27.
16. Blom, P.M.; Bakkes, S.; Spronck, P. Modeling and adjusting in-game difficulty based on facial expression analysis. *Entertain. Comput.* **2019**, *31*, 100307.
17. Hossan, M.M.; Fouda, M.M.; Eishita, F.Z. Adaptive Game Design Using Machine Learning Techniques: A Survey. In *Proceedings of the 2024 IEEE International Conference on Internet of Things and Intelligence Systems (IoTaIS)*; IEEE: New York, NY, USA, 2024; pp. 144–150.
18. Liu, P.; Tateo, D.; Bou-Ammar, H.; Peters, J. Efficient and reactive planning for high speed robot air hockey. In *Proceedings of the 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*; IEEE: New York, NY, USA, 2021; pp. 586–593.
19. Bara, M.; Gollamudi, S.; Williams, S. Developing an Air Hockey Game in LabVIEW. In *Proceedings of the 2017 25th International Conference on Systems Engineering (ICSEng)*; IEEE: New York, NY, USA, 2017; pp. 333–336.
20. Delgado-Mata, C.; Ibáñez, J. Air Hockey iOS Game That Uses Fuzzy-Logic for Game-Balancing. In *Proceedings of the 2012 Conference on Technologies and Applications of Artificial Intelligence*; IEEE: New York, NY, USA, 2012; pp. 272–277.
21. Akşahin, B.F. Dynamic difficulty adjustment by changing enemy behavior using reinforcement learning. Master's Thesis, Istanbul Technical University, Istanbul, Türkiye, 2024.
22. Romero-Mendez, E.A.; Santana-Mancilla, P.C.; Garcia-Ruiz, M.; Montesinos-López, O.A.; Anido-Rifón, L.E. The use of deep learning to improve player engagement in a video game through a dynamic difficulty adjustment based on skills classification. *Appl. Sci.* **2023**, *13*, 8249.
23. Xiong, H. Dynamic Difficulty Adjustment: Developing an adaptive game AI with machine learning. Ph.D. Thesis, California State University, Northridge, CA, USA, 2019.
24. Fuchs, R.; Gieseke, R.; Dockhorn, A. Personalized Dynamic Difficulty Adjustment Imitation Learning Meets Reinforcement Learning. In *Proceedings of the 2024 IEEE Conference on Games (CoG)*; IEEE: New York, NY, USA, 2024; pp. 1–2.
25. Dziedzic, D. Dynamic difficulty adjustment systems for various game genres. *Homo Ludens* **2016**, *9*, 35–51.
26. Engelstätter, B.; Ward, M.R. Video games become more mainstream. *Entertain. Comput.* **2022**, *42*, 100494.
27. Bowman, N.D.; Daneels, R.; Possler, D. Excited for Eudaimonia? An Emergent Thematic Analysis of Player Expectations of Upcoming Video Games. *Psychol. Pop. Media* **2024**, *13*, 416.