IJARSCT

International Journal of Advanced Research in Science, Communication and Technology

International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal

ISSN: 2581-9429

Volume 6, Issue 10, June 2026

Impact Factor: 8.2

Autonomous Driving Intelligence: Deep Q-Learning in a Simulated Racing Environment

Prof. Tejaswini Mali, Mayuri Vitthal Barmade, Samruddhi Vilas Jadhav, Bhumika Deepak Shinde

Department of Artificial Intelligence and Data Science

ISBM College of Engineering, Pune, Nande, Maharashtra

Abstract: *Autonomous driving has emerged as one of the most promising applications of Artificial Intelligence (AI) and Reinforcement Learning (RL). This project presents an intelligent autonomous driving system that uses Deep Q-Learning (DQN) to enable a virtual race car to learn driving behavior in a simulated racing environment. The system is developed using Python, Pygame, and TensorFlow/Keras, where the agent interacts with the environment through virtual sensors and learns optimal actions using a reward-based learning mechanism. Performance is evaluated using metrics such as cumulative reward, collision rate, and learning convergence. The proposed system demonstrates the effectiveness of Deep Reinforcement Learning in autonomous navigation and serves as a strong foundation for future research on intelligent transportation systems, autonomous vehicles, and AI-driven decision-making in dynamic environments.*

Keywords: Artificial Intelligence, Autonomous Driving, Deep Q-Learning (DQN), Reinforcement Learning, Deep Neural Network, Pygame, TensorFlow, Experience Replay, Target Network, Racing Simulation, Intelligent Agent, Autonomous Navigation

I. INTRODUCTION

Artificial Intelligence (AI) and Machine Learning (ML) have revolutionized modern transportation by enabling intelligent systems capable of making decisions with minimal human intervention. Among the most advanced applications of AI is autonomous driving, where vehicles can perceive their surroundings, interpret environmental conditions, and perform driving tasks without continuous human control. Recent developments in deep learning, computer vision, and sensor technologies have accelerated research in autonomous vehicles, making them more reliable and efficient for real-world applications [1], [2]. The integration of intelligent algorithms into transportation systems has the potential to improve road safety, reduce traffic congestion, minimize fuel consumption, and significantly decrease accidents caused by human error. Consequently, autonomous driving has become one of the most active research areas in Artificial Intelligence and Intelligent Transportation Systems (ITS) [3].

Conventional autonomous driving systems primarily depend on rule-based programming or supervised learning techniques that require extensive manually labeled datasets. Although these methods perform well under predefined conditions, they often fail to adapt to dynamic and uncertain environments encountered during real-world driving [4]. Reinforcement Learning (RL) provides a promising alternative by allowing an autonomous agent to learn optimal driving behavior through continuous interaction with its environment instead of relying on predefined rules. The agent receives rewards for desirable actions and penalties for undesirable behaviors, gradually improving its driving strategy through trial and error [5], [6]. Among various RL algorithms, Deep Q-Learning (DQN) combines Q-Learning with Deep Neural Networks, enabling the agent to process high-dimensional input data and make intelligent sequential decisions efficiently [7].

The proposed project, "Autonomous Driving Intelligence: Deep Q-Learning in a Simulated Racing Environment," focuses on developing an intelligent autonomous driving agent capable of learning efficient navigation strategies within a virtual racing environment. The simulation is developed using Python, TensorFlow/Keras, and Pygame, providing a safe and cost-effective platform for training and evaluating autonomous driving algorithms **[8]**. The driving agent continuously observes the racing environment using virtual sensors that provide information such as vehicle position, orientation, speed, and distances from track boundaries. Based on these observations, the Deep Q-Network predicts the optimal driving action, including steering, acceleration, braking, or maintaining speed. The learning process is further enhanced using Experience Replay and Target Networks, which improve training stability, accelerate convergence, and prevent instability during reinforcement learning [9], [10].

Simulation-based training has become an essential approach in autonomous vehicle research because real-world experimentation involves significant financial costs, safety risks, and regulatory constraints [11]. A simulated environment enables researchers to repeatedly evaluate different driving scenarios, reward mechanisms, and learning strategies without causing physical damage or endangering human lives. The proposed system demonstrates how Deep Reinforcement Learning can successfully learn adaptive driving behavior through continuous interaction with the environment while improving navigation accuracy and collision avoidance [12], [13]. The project provides a strong foundation for future intelligent transportation systems and can be further extended by integrating advanced reinforcement learning algorithms, realistic traffic conditions, multi-agent environments, and real-time sensor technologies for practical autonomous driving applications [14], [15].

II. PROBLEM STATEMENT

The development of autonomous driving systems requires intelligent decision-making in dynamic and uncertain environments, where conventional rule-based and supervised learning approaches often fail to adapt to changing road conditions and unforeseen obstacles [4], [5]. Real-world testing of autonomous vehicles is expensive, time-consuming, and involves significant safety and regulatory challenges [11]. Therefore, there is a need for an intelligent learning framework that enables an autonomous agent to perceive its surroundings, make optimal driving decisions, avoid collisions, and continuously improve its performance through interaction with the environment. This project addresses these challenges by developing an autonomous driving agent based on Deep Q-Learning (DQN) in a simulated racing environment, where the agent learns efficient navigation strategies using a reward-based learning mechanism, virtual sensor inputs, Experience Replay, and Target Networks to achieve safe, adaptive, and stable autonomous driving behavior [7], [9], [10].

III. OBJECTIVES

1. To develop an intelligent autonomous driving system using the Deep Q-Learning (DQN) algorithm in a simulated racing environment.
2. To design a realistic racing simulation with virtual sensors that enable the autonomous agent to perceive its surroundings and navigate safely.
3. To implement an effective reward mechanism, Experience Replay, and Target Network to improve learning stability and optimize driving performance.
4. To train the autonomous agent to make optimal driving decisions such as steering, acceleration, and braking while minimizing collisions and off-track movements.
5. To evaluate the performance of the proposed system using metrics such as cumulative reward, collision rate, learning convergence, and navigation efficiency.
6. To reduce human involvement in hazardous environments.

IV. LITERATURE SURVEY

1. A Novel Approach to Autonomous Driving Using Double Deep Q-Network (DDQN) in CARLA Simulation (2025)

Authors: A. Khlifi et al.

Journal: World Electric Vehicle Journal (MDPI), 2025

This paper proposes a Double Deep Q-Network (DDQN) framework for autonomous driving in the CARLA simulation environment. The model improves driving performance by reducing the overestimation problem associated with traditional DQN algorithms. The system learns lane keeping, obstacle avoidance, and safe navigation through reinforcement learning. Experimental results show that DDQN achieves higher driving stability and better decision-making accuracy than conventional DQN, making it more suitable for complex autonomous driving scenarios. However, the study mainly focuses on simulated environments and does not address real-world deployment challenges.

2. Exploring Deep Q-Network for Autonomous Driving Simulation in Different Driving Modes (2024)

Authors: M. A. Setiawan et al.

Journal: Future Artificial Intelligence and Technology for Humanity (FAITH), 2024

The authors implemented a Deep Q-Network (DQN) model capable of learning autonomous driving behavior under safe, normal, and aggressive driving modes. The reinforcement learning agent adapts its driving policy based on different environmental conditions and reward structures. The results demonstrate improved adaptability and decision-making performance across multiple driving styles. However, the model is evaluated only in simulation and requires further validation under real traffic environments.

3. Deep Reinforcement Learning for Autonomous Driving with Highway and Parking Scenarios (2024)

Authors: Y. Wang et al.

Conference: ACM International Conference, 2024

This research presents a Deep Reinforcement Learning framework for autonomous vehicles operating in highway merging and parking environments. The proposed method enables intelligent decision-making through continuous interaction with the environment and optimized reward functions. Experimental evaluation demonstrates improved navigation efficiency, collision avoidance, and policy optimization compared to conventional reinforcement learning approaches. The authors highlight that computational complexity and extensive training time remain major limitations.

4. Deep Reinforcement Learning for Autonomous Vehicles (2025)

Author: Srinivasa Kalyan Vangibhurathachhi

Journal: International Journal of Innovative Research and Creative Technology (IJIRCT), 2025

This study investigates the application of Deep Reinforcement Learning (DRL) to enhance autonomous vehicle decision-making and driving performance. The research discusses how DRL algorithms improve lane following, adaptive navigation, obstacle avoidance, and real-time control in dynamic environments. The study concludes that reinforcement learning significantly improves autonomous vehicle intelligence but requires more robust safety mechanisms and efficient training strategies for practical deployment.

5. Reinforcement Learning Methods for Autonomous Driving: A Survey (2025)

Author: Yujie Lin

Journal: Applied and Computational Engineering, 2025

This survey provides a comprehensive review of recent reinforcement learning techniques used in autonomous driving, including DQN, DDQN, PPO, DDPG, and SAC. The paper compares different reinforcement learning algorithms based on learning efficiency, decision-making capability, safety, and computational performance. It identifies major research challenges such as sparse rewards, sample efficiency, sim-to-real transfer, and safety-critical decision-making while suggesting future directions for intelligent autonomous driving systems.

IJARSCT

International Journal of Advanced Research in Science, Communication and Technology

International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal

ISSN: 2581-9429

Volume 6, Issue 10, June 2026

Impact Factor: 8.2

V. WORKING OF SYSTEM

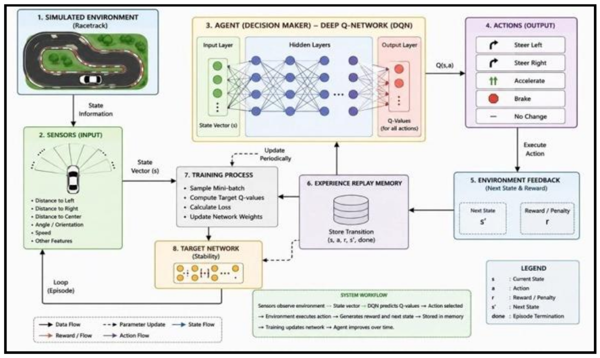

Fig 1: System architecture

A. Simulated Environment and Sensor Input

The first stage of the system consists of the **Simulated Environment** and **Sensor Input** modules. The simulated racetrack acts as the learning environment where the autonomous vehicle operates. It continuously provides information about the vehicle's position, track boundaries, obstacles, and movement. The sensor module collects essential driving information such as the distance to the left, right, and center of the track, the vehicle's orientation angle, speed, and other environmental features. These sensor readings are converted into a numerical **state vector (S)**, which accurately represents the current condition of the vehicle and its surroundings. This state vector serves as the primary input for the Deep Q-Network (DQN), allowing the autonomous agent to understand the environment before making any driving decision.

B. Deep Q-Network (DQN) Agent and Decision Making

The Deep Q-Network (DQN) functions as the intelligent decision-making component of the autonomous driving system. The state vector received from the sensor module is processed through multiple hidden layers of a deep neural network. The network analyzes the input features and predicts Q-values, which represent the expected future reward for every possible driving action. Based on these predicted Q-values, the agent selects the action that is expected to produce the maximum cumulative reward. The available actions include steering left, steering right, accelerating, braking, or maintaining the current state (no change). During training, the DQN uses an ε-greedy exploration strategy, enabling the agent to explore different actions initially and gradually exploit the learned optimal driving policy as training progresses.

C. Action Execution and Environment Feedback

Once the optimal action is selected by the DQN, it is executed in the simulated racing environment. The vehicle performs the selected driving command, such as turning left, turning right, accelerating, braking, or maintaining its speed. After executing the action, the environment immediately generates feedback in the form of a new state (S') and a reward or penalty (R). Positive rewards are assigned when the vehicle follows the track correctly, maintains appropriate speed, and successfully progresses toward the finish line. Negative rewards or penalties are given when the vehicle collides with track boundaries, moves off the track, or performs unsafe driving actions. This feedback allows the reinforcement learning agent to evaluate the effectiveness of its decisions and improve future driving performance.

D. Experience Replay Memory and Training Process

The experience generated after each interaction with the environment is stored inside the Experience Replay Memory as a tuple consisting of (Current State, Action, Reward, Next State, Done). Instead of learning only from the latest experience, the system randomly samples mini-batches of previous experiences from this memory during training. This process reduces the correlation between consecutive experiences, improves learning stability, and increases data efficiency. The sampled experiences are then used to compute target Q-values, calculate the loss function, and update the neural network weights through backpropagation. As the training process continues over multiple episodes, the

DQN gradually learns more accurate action-value functions, resulting in improved driving behavior and faster convergence.

E. Target Network and Continuous Learning Workflow

To further stabilize the learning process, the system employs a separate Target Network, which is periodically updated using the weights of the main Deep Q-Network. The Target Network generates stable target Q-values during training, preventing large fluctuations and improving convergence. Throughout the entire workflow, the sensors continuously observe the environment, the DQN predicts optimal actions, the environment provides rewards and next states, experiences are stored in replay memory, and the neural network is repeatedly updated. This cyclic learning process enables the autonomous driving agent to continuously improve its navigation strategy, reduce collision frequency, maintain lane stability, and maximize cumulative rewards. Over successive training episodes, the system evolves from making random driving decisions to performing smooth, intelligent, and autonomous navigation within the simulated racing environment.

VI. SYSTEM DESIGN

1. The proposed **Autonomous Driving Intelligence** system is designed using a **Deep Q-Learning (DQN)** architecture that enables an autonomous vehicle to learn efficient driving behavior within a simulated racing environment. The system follows the principles of **Reinforcement Learning (RL)**, where an intelligent agent continuously interacts with the environment, observes its current state, performs suitable driving actions, and learns from the rewards or penalties received. The complete system consists of several interconnected modules, including the simulated racing environment, sensor input system, Deep Q-Network (DQN), action execution module, environment feedback mechanism, experience replay memory, training process, and target network. These components work together to achieve adaptive and intelligent autonomous navigation.
2. The **simulated racing environment** serves as the training platform where the autonomous vehicle operates. It provides a realistic virtual racetrack containing curves, boundaries, checkpoints, and obstacles. The environment continuously updates the vehicle's position, orientation, speed, and collision status based on the actions selected by the agent. A sensor input module captures important environmental information such as the distance to the left, right, and center of the track, vehicle orientation, speed, and other relevant features. These sensor readings are converted into a numerical state vector, which accurately represents the current condition of the vehicle and serves as the input to the Deep Q-Network.
3. The Deep Q-Network (DQN) functions as the decision-making engine of the system. It receives the state vector from the sensor module and processes it through multiple hidden layers of a deep neural network to estimate the Q-values associated with each possible driving action. Based on these Q-values, the agent selects the optimal action, such as steering left, steering right, accelerating, braking, or maintaining its current speed. During training, the DQN follows an ε-greedy **exploration strategy**, allowing the agent to balance exploration of new actions with exploitation of previously learned driving policies. This approach enables the vehicle to gradually improve its driving performance through continuous interaction with the environment.
4. After executing the selected action, the environment generates **feedback** consisting of the next state and a corresponding reward or penalty. Positive rewards are assigned when the vehicle successfully follows the racing track, maintains stable movement, and avoids collisions, while penalties are imposed for unsafe behaviors such as crashing or leaving the track. Every interaction is stored in an **Experience Replay Memory** as a tuple containing the current state, action, reward, next state, and episode status. During training, random mini-batches are sampled from this replay memory to reduce correlation between consecutive experiences and improve learning stability. A separate **Target Network** is periodically updated to generate stable target Q-values, reducing oscillations during training and accelerating convergence.
5. The overall system operates in a continuous learning cycle where the sensors observe the environment, the Deep Q-Network predicts the best driving action, the environment executes the selected action, and the feedback is used

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 10, June 2026
Impact Factor: 8.2

to update the neural network. This iterative process enables the autonomous agent to progressively learn optimal driving strategies, minimize collisions, improve navigation accuracy, and maximize cumulative rewards. The modular architecture of the proposed system provides flexibility, scalability, and efficient training, making it suitable for future enhancements such as multi-agent learning, advanced reinforcement learning algorithms, real-time sensor integration, and deployment in more realistic autonomous driving environments.

VII. RESULTS

The proposed Autonomous Driving Intelligence system was successfully implemented and evaluated in a simulated racing environment using the Deep Q-Learning (DQN) algorithm. During execution, the autonomous agent navigated the racing track by continuously sensing the environment, selecting appropriate driving actions, and learning from reward-based feedback. The simulation demonstrates that the trained agent can follow the racing lane, negotiate curves, and avoid collisions while maintaining a stable driving speed. The on-screen performance indicators show that the vehicle successfully accumulated 20 reward points while maintaining a speed of 15 units, indicating effective learning and smooth navigation within the environment.

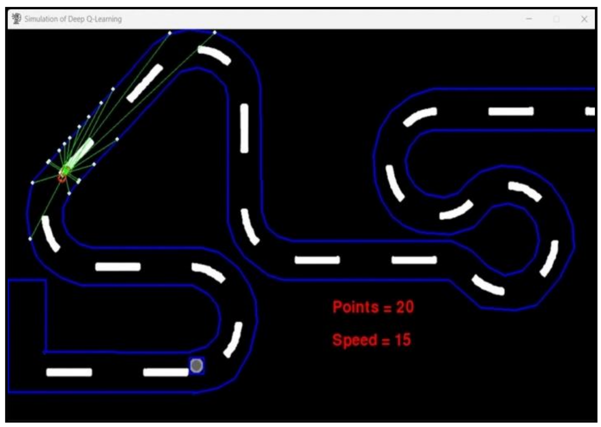

Fig 2: Output of Project

A. Simulation Performance Analysis

The simulation results indicate that the autonomous agent effectively utilized sensor information to determine optimal steering and acceleration commands. The green sensor rays continuously detected the distance to track boundaries, enabling the agent to remain within the racing lane even during sharp turns. As training progressed, the Deep Q-Network improved its driving policy by maximizing cumulative rewards and minimizing penalties due to collisions or off-track movements. The stable movement of the vehicle throughout the simulation demonstrates that the reinforcement learning model successfully learned adaptive driving behavior in a dynamic environment.

B. Performance Evaluation

The trained model achieved smooth vehicle navigation with consistent speed and efficient lane following throughout the simulated racing track. The reward mechanism encouraged the vehicle to move forward while discouraging unsafe driving behavior. Experience Replay and Target Network techniques contributed to stable convergence during training by reducing fluctuations in Q-value estimation. The simulation confirms that the proposed Deep Q-Learning approach can successfully perform autonomous navigation, making it suitable for further research in intelligent transportation systems and autonomous vehicle development.

C. Experimental Results Table

Parameter	Observed Result
Algorithm Used	Deep Q-Learning (DQN)
Simulation Environment	Python Pygame Racing Track

IJARSCT
ISSN: 2581-9429
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
Volume 6, Issue 10, June 2026
Impact Factor: 8.2

State Representation	Virtual Sensor Inputs
Available Actions	Left, Right, Accelerate, Brake, No Action
Reward Points Achieved	**20**
Vehicle Speed	**15 Units**
Lane Following	Successful
Collision Avoidance	Successfully Maintained
Learning Method	Reinforcement Learning
Overall Performance	Stable and Efficient Autonomous Navigation

VII. CONCLUSION

The proposed Autonomous Driving Intelligence using Deep Q-Learning in a Simulated Racing Environment successfully demonstrates the application of Reinforcement Learning for autonomous vehicle navigation. The Deep Q-Network (DQN) enables the intelligent agent to learn optimal driving behavior through continuous interaction with the simulated environment without requiring predefined driving rules. By utilizing virtual sensor inputs, reward-based learning, Experience Replay, and a Target Network, the agent gradually improves its ability to follow the racing track, avoid collisions, and make accurate driving decisions. The simulation results show stable navigation, efficient lane following, and consistent vehicle control, confirming the effectiveness of Deep Reinforcement Learning for autonomous driving applications. The proposed system provides a safe, cost-effective, and scalable platform for developing and evaluating intelligent driving algorithms before their deployment in real-world environments.

VIII. FUTURE SCOPE

- Integrate advanced reinforcement learning algorithms such as Double DQN (DDQN), Dueling DQN, PPO, or Soft Actor-Critic (SAC) to improve learning efficiency and driving performance.
- Extend the simulation to include realistic traffic conditions, pedestrians, traffic signals, weather variations, and dynamic obstacles for more practical autonomous driving scenarios.
- Incorporate computer vision techniques using camera images, LiDAR, and radar sensors to enhance environmental perception and obstacle detection.
- Implement multi-agent reinforcement learning to enable interaction and cooperation among multiple autonomous vehicles in complex traffic environments.
- Deploy the trained model on autonomous robotic platforms or connected vehicles using real-time embedded hardware for practical testing and intelligent transportation applications.

REFERENCES

1. Mnih, V., Kavukcuoglu, K., Silver, D., et al., “Human-Level Control Through Deep Reinforcement Learning,” *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.
2. Sutton, R. S., and Barto, A. G., *Reinforcement Learning: An Introduction*, 2nd Edition, MIT Press, 2018.
3. Kiran, B. R., Sobh, I., Talpaert, V., et al., “Deep Reinforcement Learning for Autonomous Driving: A Survey,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 6, pp. 4909–4926, 2022.
4. Dosovitskiy, A., Ros, G., Codevilla, F., Lopez, A., and Koltun, V., “CARLA: An Open Urban Driving Simulator,” *Proceedings of the 1st Annual Conference on Robot Learning (CoRL)*, pp. 1–16, 2017.
5. Khlifi, A., et al., “A Novel Approach to Autonomous Driving Using Double Deep Q-Network-Based Deep Reinforcement Learning,” *World Electric Vehicle Journal*, vol. 16, no. 3, 2025.
6. Park, Y., et al., “A Comparative Study of Deep Reinforcement Learning Algorithms for Urban Autonomous Driving Addressing the Geographic and Regulatory Challenges in CARLA,” *Applied Sciences*, vol. 15, no. 12, 2025.

7. Fei, Y., Xing, L., Yao, L., Yang, Z., and Zhang, Y., "Deep Reinforcement Learning for Decision Making of Autonomous Vehicle in Non-Lane-Based Traffic Environments," *PLOS ONE*, vol. 20, no. 4, 2025.
8. Uppuluri, B., Patel, A., Mehta, N., Kamath, S., and Chakraborty, P., "CuRLA: Curriculum Learning Based Deep Reinforcement Learning for Autonomous Driving," *arXiv preprint arXiv:2501.04982*, 2025.
9. Li, Z., et al., "A Survey of Reinforcement Learning-Based Motion Planning for Autonomous Driving," *arXiv preprint arXiv:2503.23650*, 2025.
10. Delavari, E., Khanzada, F. K., and Kwon, J., "A Comprehensive Review of Reinforcement Learning for Autonomous Driving in the CARLA Simulator," *arXiv preprint arXiv:2509.08221*, 2025.
11. Hossain, J., "Autonomous Driving with Deep Reinforcement Learning in CARLA Simulation," *arXiv preprint arXiv:2306.11217*, 2023.
12. Lillicrap, T. P., et al., "Continuous Control with Deep Reinforcement Learning," *International Conference on Learning Representations (ICLR)*, 2016.
13. Fujimoto, S., Van Hoof, H., and Meger, D., "Addressing Function Approximation Error in Actor-Critic Methods," *International Conference on Machine Learning (ICML)*, 2018.
14. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O., "Proximal Policy Optimization Algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
15. Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S., "Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor," *International Conference on Machine Learning (ICML)*, 2018.
16. Van Hasselt, H., Guez, A., and Silver, D., "Deep Reinforcement Learning with Double Q-Learning," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 30, no. 1, 2016.
17. Wang, Z., Schaul, T., Hessel, M., Hasselt, H. V., Lanctot, M., and Freitas, N., "Dueling Network Architectures for Deep Reinforcement Learning," *International Conference on Machine Learning (ICML)*, 2016.
18. Schaul, T., Quan, J., Antonoglou, I., and Silver, D., "Prioritized Experience Replay," *International Conference on Learning Representations (ICLR)*, 2016.
19. Silver, D., Huang, A., Maddison, C. J., et al., "Mastering the Game of Go with Deep Neural Networks and Tree Search," *Nature*, vol. 529, no. 7587, pp. 484–489, 2016.
20. Goodfellow, I., Bengio, Y., and Courville, A., *Deep Learning*, MIT Press, Cambridge, MA, 2016