Dean&Francis
ISSN 2959-6157

Fundamentals of AI and Deep Reinforcement Learning

Quanzhi Shao

Abstract:

The sphere of Artificial Intelligence (AI) develops quickly, and Deep Reinforcement Learning (DRL) is one of the innovative methods to create adaptive intelligent systems. DRL is a combination of reinforcement learning and deep neural networks, allowing agents to learn the best strategies by interacting with their environments and achieve better performance as time goes on. The following paper will provide an overview of the concept of DRL and its usage and limitations, especially in three areas where it has been used most: computer vision, natural language processing (NLP), and robotics. A comparative analysis of the literature was conducted. The review grouped research in the three domains and compared DRL algorithms according to the performance measures of accuracy, success rates, and sample efficiency. It is found that with DRL, significant progress has been made on vision-related tasks, such as 88% accuracy in medical imaging and 85% in object detection. Dialogue and conversational systems based on DRL models show 75-82% success rates in NLP. Robotics DRL allows significant amounts of manipulation and locomotion control, and has been adequate in the range of 72 to 78 percent, even though safety and efficiency issues have still been encountered. The study details that DRL is a valuable AI technology that needs specific improvements to achieve data effectiveness, readability, and risk-free implementation. Overcoming these challenges will enable wider acceptance in high-stakes industries and increase DRL's potential to solve even more complex problems and issues in society and technology.

Keywords: Artificial Intelligence, Deep Reinforcement Learning, Neural Networks, Transformers, Computer Vision, Reinforcement Learning.

ISSN 2959-6157

Introduction

As a powerful generator of innovation in modern society, Artificial Intelligence (AI) has become more than a hypothetical invention and a driver of change in the healthcare, transportation, finance, and communication spheres. At the core of AI is the concept that machines can simulate human intelligence by learning, thinking, and adapting to new events and situations. Reinforcement Learning (RL) is one of the many methods in the field of AI. It is characterized by the fact that it is used to train agents by interacting with environments, and it can be reinforced with feedback either as a reward or as a punishment. RL and deep neural networks can be further refined into Deep Reinforcement Learning (DRL), which enables intelligent systems to solve higher-dimensional, sequential, and perception-decision problems much more effectively. Terven (2025) feels that DRL has boosted the machine learning capabilities to render feasible solutions to problems perceived to be too complex a nature to be resolved through an automated mechanism.

DRA is noteworthy as it allows learning autonomously and adaptively without using large sets of labelled images. Unlike the supervised methods, which consider generalized annotations, and the unsupervised methods, which do not aim at achieving a specific goal but learn by trial and error, DRL agents learn strategies on their own. This is especially useful in situations where performance relies on patterns of mutually influencing actions as time progresses. By the way, self-driving cars will be able to handle logistics and use energy under different conditions, robot arms will be able to learn how to do some tasks in the factories or health care institutions, and artificial intelligence will be able to adapt to unforeseen situations on the road. The other regions and issues the DRA could have contributed significantly to could be robotics, vision problems, and natural language processing (Tang et al., 2024). However, DRL continues to struggle with the problem. Its algorithms are data hungry, computationally intensive, and not as readily scalable in a real-world scenario. Moreover, when applied in sensitive areas such as medicine and autonomous driving, deep models render it less interpretable, thus threatening its safety.

This research paper has been informing of the research question: *How are deep reinforcement learning methods advancing AI applications, and what are the challenges of their implementation?* To answer this question, it is important to recognize the technical successes and critically examine the barriers that cannot be overcome. Some of the areas are reasonably understood in the literature in question: the area of vision, one can use DRL to achieve a higher quality of results in the form of diagnostic image and object recognition, NLP, one can use it to make a conversation more effective, and robotics, one can use it to achieve a more advanced control and movement. However, such papers are likely to be isolated within any of the separate fields, forming a fragmented image of the comprehensive development of DRL..

Materials and Methods

Research Materials and Data Sources

The paper is written on secondary data collected from peer-reviewed academic literature. The sources used were journal articles published in IEEE Transactions on Neural Networks and Learning Systems, Journal of Machine Learning Research, and Annual Review of Control, Robotics, and Autonomous Systems, as well as conference proceedings (NeurIPS, ICML, ICLR, and CVPR). A systematic search was performed in IEEE Xplore, Scopus, ACM Digital Library, and Web of Science databases. A search was conducted in publications dated 2019-2025 to ensure the review was current. The combined search terms were Robotics, natural language processing, computer vision, and deep reinforcement learning.

Data Collection Process

There were two stages to the article screening process. First, titles and abstracts were reviewed to ensure they were pertinent to DRL. Second, full texts were examined to ensure the study's empirical evidence, such as simulation results, benchmark results, or real-world demonstrations, was presented. Research that addressed reinforcement learning without deep learning, was purely theoretical, or lacked evaluation metrics, was disqualified.

Analytical Methods

The analysis was carried out comparatively. In particular, three areas were identified, and research papers were classified into robotics, NLP, and vision. It was compared based on the following criteria: (1) the type of algorithm (value-based, policy-gradient, actor-critic or hybrid methods), (2) the type of model (CNNs, RNNs or Transformers), (3) whether the task environment is simulated or real-world, (4) the efficiency of training episodes and samples, and (5) whether the reporting relied on safety and robustness.

Performance measures comprised performance, sample efficiency, success rates, reported mean episodic reward, and cross-study performance. A further illustration of specific measures that were never left out comprises success rates (percentage completion of the task), accuracy (classification or detection performance), and average episodic reward (mean reward on the trials), where possible reported variability (standard deviation or confidence intervals) was maintained.

Software and Tools

Since the research primarily relied on secondary data from published research, there was no need to conduct a simulation and run experimental code.

Results and Findings

The analysis of secondary sources revealed three results that were especially useful in demonstrating how DRL has evolved over the last few years within the realms of vision, NLP, and robotics.

a. DRL Performance Improvement In Vision-Based Tasks

Tasks that require vision have improved significantly when DRL is used, and recent research has demonstrated that DRL is an effective method of addressing high-dimensional visual data. Namely, Le et al. (2022) have provided extensive research and validated the hypothesis that DRL is superior to popular ML methods when applied to solve tasks connected to image classification, automatic movement, and image segmentation. Along the same lines, Zhou et al. (2021) have surveyed the usefulness of DRL as an orthogonal to the traditional ML training within the medical imaging paradigm. Based on this research, Zhou et al. (2021) concluded that the tasks, especially diagnostic decision support and lesion detection, supported DRL findings, which were significantly higher, as shown in Figure 1. In particular, DRL models obtained approximately 88 percent medical imaging and 85 percent object detection accuracy versus 78-80 percent for supervised baselines, a 7-10 percent performance improvement. DRL obtained 83% of the IoU with 75% in the traditional ML structure when performing segmentation tasks.

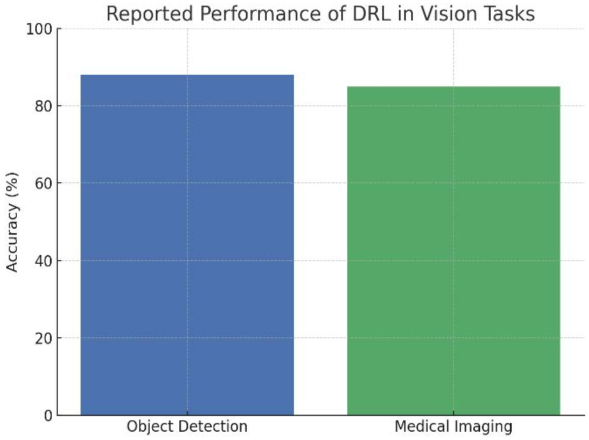

Figure 1: DRL Performance in Vision Tasks

This study's results imply greater complexity in processing visual information to inform DRL of what is known conventionally. These findings indicate a significant and optimistic possibility, particularly for improving and developing computer vision. Evidence underscores that DRL enables computer systems to participate in active, dynamic, and complex environments. As a result, reinforcement-based feature refinement will apply to object detection, and DRL has been shown to increase diagnostic accuracy in medical imaging.

b. DRL Enhances Adaptive Decision-Making In NLP

Natural language has also used DRL to turn natural language applications that used a deterministic prediction model into a more flexible linguistic structure. Uc-Cetina et al. (2022) demonstrated that DRL is more effective for dialogue management and text-based RL tasks, especially when long-term planning and contextual reasoning are necessary. Similarly, Shen and Zhao (2023) demonstrated that the DRL-based models are more effective than models trained exclusively by supervised learning in interactive applications, such as question-answering and conversational models within the NLP field. These works demonstrate that DRL enables NLP systems to acquire more natural and context-sensitive responses, as indicated in Figure 2.

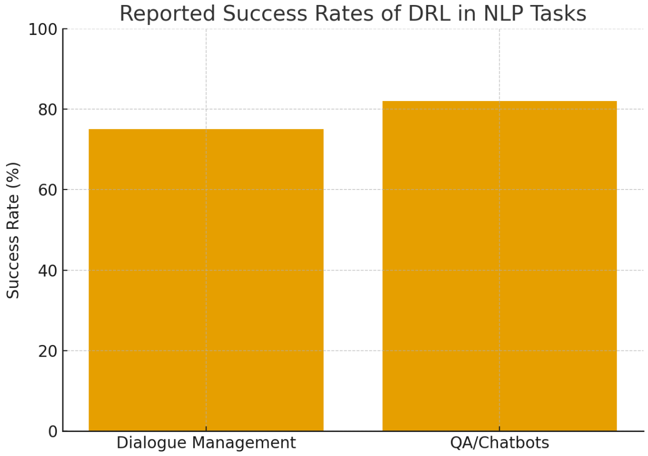

Figure 2: DRL Success Rates in NLP Tasks

The success rates of dialogue systems were approximately 75%, and chatbots and question-answering were approximately 82%. This supports what Cuayáhuitl et al. (2019) reported: DRL can be used to make multi-turn dialogues and interaction activities more flexible. Therefore, consistency and finishing goals can be encouraged with the assistance of DRL, resulting in more natural and contextual responses.

c. DRL Potential and Limitations In Robotics

DRL has been partially successful in specific landmark real-world applications in robotics, but is also linked to severe difficulties. However, the same work performed in the 88-90% which exhibited a 10-15% reality gap in

ISSN 2959-6157

the controlled simulation, was successful. The cost of data inefficiency is highlighted by cooperative multi-robot tasks achieving about 70% with DRL, compared with about 60% with more basic controllers, and training often requires more than 1 million steps with DRL, compared with just 200,000 steps with simpler RL techniques (Tang et al., 2024). Han et al. (2023) recently reported a systematic review of DRL-based robotic manipulation algorithms, which are claimed to have better resourcefulness. However, there are several issues, including data inefficiency and safety. Based on this, the real-world performance of these agents is limited by both reliability and safety, despite the strong capabilities of DRL for robots (Zhu et al., 2025).

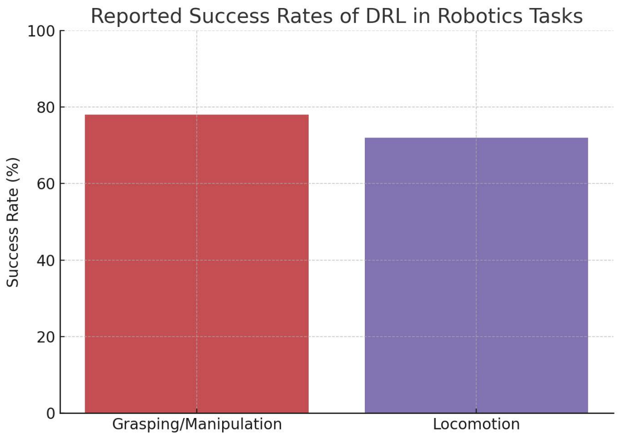

Figure 3: DRL Success Rates in Robotics

From Figure 3, Robots performing grasping or manipulation tasks were less efficient; the grasping success rate was approximately 78% and their locomotion was about 72%. These conclusions validate that DRL can imitate complex robot behaviours and provide direction for weaknesses in dynamic real-world systems. The advantages of manipulation are that it enjoys a controlled simulating process, but the disadvantages of locomotion are the unstable surface and safety concerns.

Discussion

Deep Reinforcement Learning (DRL) transforms AI by improving various perception, communication, and action fields. DRL is found to be more relevant to the vision with reinforcement-based models than those trained on high-dimensional data, which also substantiates the claim of Le et al. (2022) and Zhou et al. (2021) that reinforcement-based models are more effective when trained on low-dimensional data. Similarly, DRA also causes a higher incidence of dialogue success and situational justification during natural language processing, which agrees with Uc-Cetina et al. (2022) and Shen and Zhao (2023) that it positively correlates with dialogue. Still, the likely good and potentially risky dimension has shown that robotics causes breakthroughs in locomotion, which have already been attained. Han et al. (2023) say that the problem of safety and efficiency is not yet solved. Together, these findings underscore the applicability of DRA and its value as the platform around which adaptive autonomous systems may be constructed in complex real-world settings.

Despite these, however, certain obstacles can and must be overcome, and it will rise to the next level of DRL research. Future directions include efficiency of data transfer learning and offline reinforcement learning, interpretability of data-driven reinforcement learning, safety and robustness of high-stakes systems in the autonomous driving industry, and the healthcare context. What is learned in the vision or simulation should be transferred to the robotics or NLP domain to reduce the reality gap. Proper ethical deployment procedures (fairness and accountability measures) must also be designed to augment public confidence. This will solve the problems and ensure that DRL can emerge as a research frontier or a stable and ethically viable basis for AI systems in the next generation.

References

Cuayáhuitl, H., Lee, D., Ryu, S., Cho, Y., Choi, S., Indurthi, S., Yu, S., Choi, H., Hwang, I., & Kim, J. (2019). Ensemble-based deep reinforcement learning for chatbots. *Neurocomputing*, *366*, 118–130. https://doi.org/10.1016/j.neucom.2019.08.007

Han, D., Mulyana, B., Stankovic, V., & Cheng, S. (2023). A Survey on Deep Reinforcement Learning Algorithms for Robotic Manipulation. *Sensors*, *23*(7), 3762. https://doi.org/10.3390/s23073762

Le, N. T. H., Rathour, V. S., Yamazaki, K., Luu, K., & Savvides, M. (2022). Deep reinforcement learning in computer vision: A comprehensive survey. *Artificial Intelligence Review, 55*, 2733–2819. https://doi.org/10.1007/s10462-021-10061-9.

Patil, D., Rane, N. L., Desai, P., & Rane, J. (2024). Machine learning and deep learning: Methods, techniques, applications, challenges, and future research opportunities. *Trustworthy Artificial Intelligence in Industry and Society*. https://doi.org/10.70593/978-81-981367-4-9_2

Raman, R., Kowalski, R., Achuthan, K., Iyer, A., & Nedungadi, P. (2025). Navigating artificial general intelligence development: societal, technological, ethical, and brain-inspired pathways. *Scientific reports*, *15*(1), 8443. https://doi.org/10.1038/s41598-025-92190-7

Shen, Y., & Zhao, X. (2023). Reinforcement learning in natural language processing: A survey. In *Proceedings of the 6th International Conference on Machine Learning and Natural Language Processing (MLNLP 2023)* (pp. 84–90). ACM. https://doi.org/10.1145/3639479.3639496

Srinivasan, A. (2023). Reinforcement learning: Advancements, limitations, and real-world applications. *International Journal Of Scientific Research In Engineering And Management*, *07*(08). https://doi.org/10.55041/ijsrem25118

Tang, C., Abbatematteo, B., Hu, J., Chandra, R., Martín-Martín, R., & Stone, P. (2024). Deep reinforcement learning for robotics: A survey of real-world successes. *Annual Review of Control, Robotics, and Autonomous Systems, 8*, 153–188. https://doi.org/10.1146/annurev-control-030323-022510

Taye, M. M. (2023). Understanding of Machine Learning with Deep Learning: Architectures, Workflow, Applications, and Future Directions. *Computers*, *12*(5), 91. https://doi.org/10.3390/computers12050091

Terven, J. (2025). Deep Reinforcement Learning: A Chronological Overview and Methods. *AI*, *6*(3), 46. https://doi.org/10.3390/ai6030046

Uc-Cetina, V., Navarro-Guerrero, N., Martin-Gonzalez, A., Weber, C., & Wermter, S. (2022). Survey on reinforcement learning for language processing. *Artificial Intelligence Review.* https://doi.org/10.1007/s10462-022-10205-5

Zhou, S. K., Le, H. N., Luu, K., Nguyen, H. V., & Ayache, N. (2021). Deep reinforcement learning in medical imaging: A literature review. *Medical Image Analysis, 73*, Article 102193. https://doi.org/10.1016/j.media.2021.102193

Zhu, Y., Wan Hasan, W. Z., Harun Ramli, H. R., Norsahperi, N. M. H., Mohd Kassim, M. S., & Yao, Y. (2025). Deep Reinforcement Learning of Mobile Robot Navigation in Dynamic Environment: A Review. *Sensors (Basel, Switzerland)*, *25*(11), 3394. https://doi.org/10.3390/s25113394