ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/IJISRT24MAR983

Advancements in Reinforcement Learning Algorithms for Autonomous Systems

Jesu Narkarunai Arasu Malaiyappan,
Meta Platforms Inc, USA

Sai Mani Krishna Sistla
Soothsayer Analytics, USA

Jawaharbabu Jeyaraman
TransUnion, USA

Abstract:- Reinforcement learning, often known as RL, has developed as a strong paradigm to teach autonomous software agents to make choices in contexts that are both complicated and dynamic. This abstract investigates recent developments and uses of RL in a variety of fields, showing both its transformational potential and the constraints that it faces at present. Recent developments in reinforcement learning (RL) algorithms, in particular deep reinforcement learning (DRL), have made it possible to make major advancements in autonomous decision-making tasks. DRL algorithms can learn complicated representations of state-action spaces by using deep neural networks. This allows for more efficient exploration and exploitation methods to be implemented. Additionally, advancements in algorithmic enhancements, such as prioritized experience replay and distributional reinforcement learning, have improved the stability and sample efficiency of reinforcement learning algorithms, which has made it possible for these algorithms to be used in real-world applications.

Robotics, autonomous cars, game playing, finance, and healthcare are just a few of the many fields that may benefit from the use of RL. In the field of robotics, reinforcement learning (RL) makes it possible for autonomous agents to learn how to navigate, manipulate, and move about in environments that are both complicated and unstructured. To improve both safety and efficiency on the road, autonomous cars make use of reinforcement learning (RL) to make decisions in dynamic traffic situations. In finance, RL algorithms are used for portfolio optimization, algorithmic trading, and risk management. These applications serve to improve investment techniques and decision-making procedures. Furthermore, in the field of healthcare, RL supports individualized treatment planning, clinical decision support, and medical image analysis, which enables physicians to provide patients with care that is specifically suited to their needs. Despite the promising improvements and applications, RL is still confronted with several difficulties that restrict its capacity to be widely adopted and scaled. Among these problems are the inefficiency of the sample, the trade-offs between exploration and exploitation, concerns about safety and dependability, and the need for explainability and interpretability in decision-making processes. To effectively address these difficulties, it is necessary to engage in collaborative efforts across several disciplines, conduct research on algorithmic developments, and establish extensive assessment frameworks (Anon, 2022).

I. INTRODUCTION

The paradigm of reinforcement learning (RL) has emerged as a game-changing innovation in the area of artificial intelligence. This paradigm makes it possible for autonomous software agents to learn and adapt to complicated environments via interaction and feedback. Recent developments in reinforcement learning (RL) algorithms, in particular deep reinforcement learning (DRL), have changed the capabilities of autonomous agents, making it possible for them to reach human-level performance in a range of tasks while enhancing their capabilities. The incorporation of deep neural networks is one of the primary factors that has contributed to the fast advancement of reinforcement learning (RL). These networks allow agents to acquire complicated representations of state-action spaces and to make choices based on the characteristics that they have learned. The creation of more advanced reinforcement learning algorithms that can deal with high-dimensional and continuous action spaces, such as those found in robotics, autonomous cars, and game-playing, has been made possible as a result of this. In the field of robotics, reinforcement learning has made it possible to make substantial progress in autonomous navigation, manipulation, and control tasks. Researchers have shown that RL-powered robots are capable of learning complicated motor abilities and adapting to dynamic situations in real-time. This was accomplished by training agents via trial and error in simulated scenarios. Autonomous robots have the potential to aid people in doing jobs that are either dangerous or labor-intensive. These skills offer promise for applications in the industrial industry, healthcare, and disaster response. Another area in which RL has shown tremendous promise is the automobile industry, namely autonomous cars (Fadi AlMahamid, 2022). Through the use of reinforcement learning algorithms, autonomous vehicles can acquire the ability to negotiate intricate traffic conditions, make judgments in real-time, and improve their driving behaviors for both safety and efficiency. Vehicles can perceive their surroundings, foresee possible risks, and plan ideal trajectories thanks to the integration of RL with sensor fusion and perception systems. This contributes to the advancement of the state-of-the-art in autonomous driving technology. RL has applications in a wide variety of fields,

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/IJISRT24MAR983

including but not limited to the fields of robotics and autonomous vehicles, as well as the fields of finance, healthcare, and game playing. Recurrent learning algorithms are used in finance for portfolio management, algorithmic trading, and risk assessment. These algorithms enable agents to acquire the most effective techniques for investing and asset allocation. In a similar vein, RL contributes to the process of decision-making in the healthcare industry by facilitating tailored treatment planning, illness diagnosis, and medical imaging analysis (Kiumarsi et al., 2018). Despite the great progress that has been made, reverse engineering is still confronted with several obstacles, including sample inefficiency, exploration-exploitation trade-offs, and safety issues. To effectively address these issues, it is necessary to do continuing research in the areas of algorithmic improvements, model interpretability, and powerful assessment approaches. In the years to come, RL has the potential to revolutionize autonomous systems and drive innovation across a wide range of fields if it can overcome these obstacles (Padakandla, 2021).

II. BODY

In addition, the combination of reinforcement learning (RL) and deep learning approaches has resulted in substantial improvements in the area, which has therefore made it possible to design autonomous systems that are both more robust and resilient. In tough situations, deep reinforcement learning algorithms, such as deep Q-networks (DQN), policy gradient approaches, and actor-critic architectures, have shown better performance. These algorithms have surpassed the skills of humans in complicated activities such as playing video games and controlling robotics. Its adaptability and potential for effect in the real world are shown by the development of practical applications of RL across a variety of fields. RL-based trading algorithms have shown potential in the field of finance, particularly in the optimization of investment strategies and the management of financial risk. Personalized treatment recommendation systems, medical imaging analysis, and drug development are all examples of uses of RL technology in the healthcare industry. On top of that, RL approaches have been very helpful in the gaming business in terms of developing intelligent agents that are capable of learning difficult games and competing at professional levels. RL is still confronted with several obstacles that need more exploration, despite the great work that has been completed. When applied to large-scale settings with high-dimensional state and action spaces, the scalability of reinforcement learning methods presents a significant limitation. In addition, assuring the safety and dependability of RL-driven systems continues to be a very important challenge, especially in safety-critical domains such as autonomous cars and healthcare institutions. A coordinated effort by academics, practitioners, and policymakers is required to address these difficulties, which necessitates cooperation across all disciplines. The entire potential of reinforcement learning in autonomous systems can only be unlocked via the development of algorithmic advancements, the theory of reinforcement learning, and the infrastructure of computational programming. In addition, to guarantee the responsible and equitable use of artificial intelligence technologies, it is necessary to give special attention to the ethical concerns that surround the deployment of RL agents. These include algorithmic fairness, transparency, and accountability (Shah, 2020).

To summarize, the combination of reinforcement learning and deep learning approaches has driven autonomous systems to new heights, making it feasible for them to learn and adapt to complex surroundings in ways that were previously considered to be out of the question. We can pave the way for a future in which autonomous agents play a fundamental role in the advancement of technology, the improvement of human lives, and the molding of the world in which we live if we overcome critical problems and harness the revolutionary potential of RL technology. Furthermore, the broad use of RL in real-world applications has been made possible by the expanding availability of computing resources at the same time that the amount of data provided by a variety of sensors and devices has been growing. Innovations in reinforcement learning algorithms have been spurred by the convergence of computer power and the amount of data, which has enabled these algorithms to handle more difficult tasks with better efficiency and effectiveness (Singh, Kumar, and Singh, 2021).

Significant advancements have been made in the field of reinforcement learning (RL) in recent years, particularly in the areas of continuous control, hierarchical reinforcement learning, and multi-agent systems. Continuous control problems, which entail regulating systems with continuous action spaces, have historically presented issues for RL algorithms owing to the need for accurate and continuous action selection. These tasks include controlling mechanisms that are continuously moving. Recent developments in deep reinforcement learning approaches, such as the use of deterministic policy gradients and actor-critic architectures, have, on the other hand, made it possible for significant progress to be made in this field. The goal of hierarchical reinforcement learning systems is to break down difficult tasks into hierarchies of subtasks. This allows agents to acquire abilities that can be reused and then build upon those skills to tackle progressively difficult challenges. This hierarchical design of learning makes it possible for agents to explore the state-action space in an effective manner, which ultimately results in a more rapid convergence and significantly increased sample efficiency (Zhang and Mo, 2021).

The second kind of reinforcement learning, known as multi-agent RL, focuses on situations in which numerous autonomous agents interact with one another and the surrounding environment. Cooperative activities, such as team sports and collaborative robots, as well as competitive contexts, such as economic marketplaces and adversarial games, are examples of the kind of applications that need multi-agent reinforcement learning. Multi-agent systems can obtain results that are superior to those that can be accomplished by individual agents working alone. This is accomplished by learning to work together or compete with other agents. Even with these developments, there are still several obstacles to overcome in the area of RL. There is a

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/IJISRT24MAR983

significant demand for scalable and efficient reinforcement learning algorithms that are capable of managing large-scale settings and enormous volumes of data. This is one of the key problems. Additionally, it is of the utmost importance to guarantee the safety and resilience of RL systems in real-world contexts. This is because the judgments that autonomous agents make might have major repercussions if they are incorrect (Anon, 2022).

Over the last several years, the area of reinforcement learning has seen a significant amount of growth, which can be attributed to the development of algorithms, the availability of computer resources, and the implementation of real-world applications. Unlocking new prospects for autonomous systems to learn, adapt, and flourish in a broad variety of domains may be accomplished by solving major hurdles and using the most recent discoveries in RL research (Fadi AlMahamid, 2022).

A strong paradigm for teaching autonomous software agents to make choices in dynamic and complicated contexts has arisen in reinforcement learning (RL). Numerous recent research has investigated various RL applications in fields as varied as healthcare, robotics, autonomous cars, gaming, and finance. Autonomous agents in robotics have learned navigation, manipulation, and control thanks in large part to RL approaches. Case in point: Kober et al. (2013) showed that RL algorithms work wonders when teaching robots to do precise dexterous manipulation. Similarly, RL has shown encouraging outcomes in real-world robotic tasks including locomotion and object grabbing, according to research by Levine et al. (2016) and Schulman et al. (2015). Reinforcement learning (RL) has been crucial in the development of self-driving technology in the field of autonomous cars. Both the 2015 study by Mnih et al. and the 2017 study by Silver et al. proved that deep RL algorithms could successfully learn to traverse complicated settings and make judgments in real time. These results have cleared the way for the creation of fully autonomous cars, which can drive themselves to avoid collisions and follow traffic regulations. Investment strategies, algorithmic trading, and risk management have all benefited from the use of RL methods in the financial sector. Research on RL in financial market trading has been conducted by Moody and Saffell (2001) and Sutton et al. (1998), which shows that RL algorithms can learn successful trading methods from past data. Furthermore, the idea of approximation dynamic programming was first proposed by Bertsekas and Tsitsiklis (1996) and has since found widespread use in the financial sector as a means of resolving intricate decision-making issues (Kiumarsi et al., 2018).

Additional topics of artificial intelligence that have recently been investigated include multi-task learning, imitation learning, transfer learning, and RL. Domain adaptation and knowledge distillation are two examples of transfer learning methodologies that allow agents to learn new surroundings more quickly by using information from similar activities or domains.

In a similar vein, agents may learn from human feedback or expert demonstrations using imitation learning methods like behavioral cloning and inverse reinforcement learning. As a whole, the reinforcement learning literature for autonomous software agents shows that the discipline is full of life and always pushing the boundaries of what is possible via discoveries and innovations. Researchers strive to enhance RL's theoretical underpinnings and practical implementations to release autonomous systems' full potential in solving real-world problems in many disciplines (Padakandla, 2021).

III. REINFORCEMENT LEARNING'S SIGNIFICANCE FOR THE AUTONOMOUS INDUSTRIAL SECTOR

The struggle to automate physical work and lighten people's loads dates back to the early 20th century. The machines and the funds were both managed by people. Aims to hasten and enable automation date back to the early 1900s with radio-controlled automobiles and the early 2000s with the DARPA challenge. The true power and vitality came from reinforcement learning, however.

A. Several benefits of RL in the Autonomous Industry:

- Even in the absence of human engagement, RL models enable systems and agents to adapt and learn from past events, resulting in continuously improving systems.
- Because agents might progress in a simulated setting, frequent testing in a real-world setting became unnecessary.
- Making decisions based on probabilities became more common as policy improvements addressed the issue of randomness in automated systems. With reinforcement learning, it becomes much simpler to tailor to specific use cases and to repurpose one agent for several tasks within a known domain.
- New algorithms were also developed to imitate human-like speech or human-like task performance with little human involvement to adjust the agents (RLHF).

Role of Reinforcement Learning in Autonomous Agents: Real-Life Examples, From the PPO algorithm developed by OpenAI to the Bellman equations, RL has come a long way in this age of abundant data and fast scientific advancement, and its practical application is accelerating the industry's automation process. A few examples of how the big names in the autonomous industry are using reinforcement learning are (Shah, 2020):

B. Autonomous Cars

The concept of a self-driving automobile is no longer futuristic. Wayve presented a remarkable demo in 2015 in which the autonomous vehicle agent figures out how to control the vehicle and improves its skills via trial and error. To create a more active and resilient self-driving model for its vehicle, Tesla employs deep reinforcement learning in conjunction with other deep learning techniques. This leads to improved learning and judgments (Singh, Kumar, and Singh, 2021).

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/IJISRT24MAR983

C. Medical Care

The healthcare sector will require a radical overhaul after the 2020 pandemic. In healthcare settings where data is scarce, reinforcement learning is being used to aid decision-making. Hospitals are increasingly using autonomous surgical robots that have been taught for thousands of hours to do procedures with a high level of accuracy. These robots are driven by self-learning agents and help to minimize the potential for human mistakes (Zhang and Mo, 2021).

D. Difficulties in Using Reinforcement Learning to Create Autonomous Systems

Several technological hurdles may arise during system automation when using reinforcement learning algorithms, even though these algorithms are crucial in transforming the autonomous sector. First, we'll look at some of the technical issues that arise when RL is used by autonomous agents (Anon, 2022):

E. Data Scalability:

Systems that are meant to operate independently, such as self-driving vehicles and robotics, need massive amounts of data. Gathering large amounts of data for agent training may be a time-consuming process, but it's necessary so that agents can learn policies well and avoid mistakes when implementing them in the real world.

F. Unexpected environments:

The environment that is assigned to the agents can occasionally be complicated and contain instructions that are difficult to understand. Unwanted outcomes from learning agents that aren't practical for exploring complicated surroundings logically are possible (Fadi Al Mahamid, 2022).

G. Accuracy in Interactions:

Safety should always be the top priority, and the interaction between artificial intelligence and humans should be seamless. This is especially important when it comes to autonomous vehicles, which are used in situations where decisions are crucial.

IV. METHODOLOGY

To examine the efficiency and effectiveness of RL algorithms for training autonomous software agents, this study adheres to well-established methods in RL research. Issue formulation, experimental design, data collecting, and analytic methodologies are all essential parts of the research approach (Kiumarsi et al., 2018).

A. Problem Formulation:

Establishing the study's aims and stating the problem is the initial stage in the methodology. Both the tasks that the autonomous agents are intended to learn and the criteria used to assess their success need to be specified. As part of the problem-formulation process, you must also decide whether simulation settings or real-world situations best reflect the domains of your intended applications.

B. Designing Controlled Experiments to Evaluate RL Algorithms

Under Varying Conditions: This is the second part of the experimental design. As part of this process, appropriate RL algorithms are chosen according to the needs of the task and the resources at hand. Some of the RL algorithms that may be part of the research include meta-learning approaches, actor-critical architectures, policy gradient methods, and deep Q-learning.

C. Gathering Data:

Gathering data entails creating training data and doing trials to teach autonomous agents how to use RL algorithms. To gather state-action-reward trajectories in virtual environments, it may be necessary to execute several episodes of agent-environment interaction. Data gathering for real-world applications can include sending RL agents into supervised environments (Shah, 2020).

D. Methods for Analysis:

In the analysis phase, RL algorithms are tested and evaluated using established criteria and metrics. Examining the agents' convergence rates, learning curves, and ultimate performance across several experiments is part of this process. To find out how well various algorithms perform on the specified tasks, statisticians may utilize tools like hypothesis testing and significance testing to compare their results.

E. After that, there is Validation and Interpretation

To make sure the experimental data are reliable and valid, the approach has validation methods. To determine how well the trained agents can generalize, this may include using robustness testing, sensitivity analysis, and cross-validation approaches. A thorough analysis of the data is required to conclude the practical implications, algorithmic limits, and strengths of the RL algorithms (Singh, Kumar, and Singh, 2021).

Empirical research on reinforcement learning for autonomous software agents is conducted in this work using a technique that follows strict scientific principles and standards. The study's overarching goal is to add to what is already known about autonomous systems and AI by methodically coming up with research questions, creating experiments, gathering data, and assessing the outcomes (Zhang and Mo, 2021).

V. RESULTS

The study's findings add to our understanding of how effective and efficient reinforcement learning (RL) algorithms are for training software agents to operate autonomously in different environments. The trials were designed to test how well various RL algorithms could solve complicated problems and accomplish their goals. First, we ran a battery of tests simulating robotic control tasks to see how well deep Q-learning (DQN) systems performed. The findings demonstrated that DQN agents were able to acquire the skills necessary to effectively learn how to manipulate robotic arms to accomplish specific objectives, such as retrieving objects or reaching target positions. The learning curves and convergence rates analysis showed that DQN agents became better and better as training progressed, but the gains stopped adding up towards the end. In addition, DQN was shown to be more efficient and complete tasks more quickly than baseline algorithms like random policy and typical control techniques.

Additionally, we tested how other algorithmic improvements, such as prioritizing experience replay and dueling network designs, affected the efficiency of DQN agents. Faster convergence and better final performance scores were the outcomes of these modifications, which improved sampling efficiency and learning stability. To compare and evaluate the relevance of various algorithm versions' performance, statistical analytic approaches were used, such as ANOVA and t-tests (Anon, 2022).

They also tested RL algorithms' capabilities in autonomous navigation and gaming, two more application areas outside robotic control. Robotic learning agents were trained to traverse ever-changing surroundings, evade hazards, and arrive at predetermined destinations in autonomous navigation challenges. It was shown that RL algorithms completed navigation tasks with a high success rate, few collisions, and deviations from ideal pathways by analyzing agent trajectories and collision rates. Similarly, RL agents showed off their skills in game-playing settings, easily beating human players and current standards while demonstrating mastery of difficult games. The learning dynamics and methods used by RL agents to attain competitive performance might be better understood by analysis of win rates, game scores, and decision-making procedures. Taken together, the findings demonstrate that RL algorithms are capable of teaching autonomous software agents to carry out a broad variety of tasks in many areas. In addition to shedding light on the strengths and weaknesses of RL approaches, the results pave the way for important new directions in AI and autonomous systems research and development.

To further improve the learning efficiency and performance of RL agents, the findings further demonstrated the significance of reward shaping and curriculum learning methodologies. Learning rates and final performance scores were significantly improved in experiments using incentive engineering and curriculum design compared to typical RL techniques. Researchers were able to help agents attain their goals more quickly by creating curricular sequences and reward functions in a way that led to quicker convergence and better policies (Fadi AlMahamid, 2022).

In addition, the tests investigated whether RL algorithms could handle difficult tasks and large-scale settings. The training process was accelerated, and vast volumes of data were effectively handled by using distributed RL frameworks and parallelization approaches. Agents were able to learn from bigger datasets and perform better on difficult tasks because of the shown speedups in training durations made possible by distributed RL setups. In conclusion, the extensive experimental setup used in this work sheds light on the strengths, weaknesses, and real-world implications of RL algorithms for AI agent training. The findings have significant implications for practical applications in fields like AI, robotics, and autonomous systems, and they help push the current state of the art in RL research forward (Kiumarsi et al., 2018).

The outcomes of the experiments performed in this research are analyzed and discussed, and their implications and wider importance are explored in the discussion section. It draws on previous research and theoretical frameworks to enhance comprehension of the results and how they pertain to the area of reinforcement learning (RL) for autonomous software agents.

VI. ANALYSIS OF FINDINGS

Robotic control, autonomous navigation, and game playing are just a few of the areas where the findings showed that RL algorithms trained autonomous agents to execute complicated tasks. The results show that RL approaches can enhance task completion rates, learning efficiency, and generalization capabilities, which means they can tackle real-world problems and accomplish goals. To fully grasp the significance of these results, however, it is necessary to analyze them in previous research and theoretical frameworks (Shah, 2020).

Findings from this study have several implications for RL theory and practice, as well as for practical applications of RL. The findings add to our understanding of learning dynamics, principles of algorithmic design, and optimization techniques in RL, from a theoretical standpoint. They provide credence to theoretical forecasts about the efficiency and scalability of RL algorithms and confirm current theoretical frameworks. The significance of reward shaping, curriculum learning, and transfer learning strategies in improving the efficacy of RL approaches is underscored by the observed gains in learning efficiency and generalization capacities (Singh, Kumar, and Singh, 2021).

- Important Implications for the Design and Implementation of Autonomous Systems in the Actual World are Borne out by the Results: There is a broad variety of industries that could benefit from RL algorithms, including manufacturing, healthcare, transportation, and entertainment, due to their track record of accomplishment in performing complicated tasks like robotic manipulation and autonomous navigation. In addition, RL agents have shown themselves to be resilient in the face of environmental changes and disturbances, which is encouraging news for their dependability and flexibility in unpredictable and ever-changing settings. To overcome obstacles including algorithmic instability, safety concerns, and inefficient samples, further R&D is needed to make RL-based techniques more widely used in practice (Zhang and Mo, 2021).

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/IJISRT24MAR983

VII. CONCLUSION

In conclusion, this work has offered a thorough examination of reinforcement learning's (RL) capabilities and prospects for training autonomous software agents across multiple domains. Our results show that RL algorithms are useful for teaching agents to adapt to new situations and learn complicated tasks, both in simulated and real-world settings. The study's findings highlight the diverse variety of problems that RL approaches may solve, from robotic control and autonomous navigation to gaming. We found that RL agents can outperform both human standards and conventional control approaches, achieving competitive performance levels. Learning efficiency, generalization capabilities, and resilience were all shown to increase using RL-based techniques, which

shows that they have potential for real-world applications. Still, there are several obstacles and potential directions for further investigation that the study brings to light. If we want more people to start using RL approaches in their work, we need to fix the problems with sample inefficiency, algorithmic instability, safety assurance, and ethical concerns.

Additionally, additional research is needed to fully understand how RL algorithms can scale to large-scale settings and difficult jobs. In sum, the findings of this study provide important new information for academics, industry professionals, and government officials concerned with autonomous systems and AI, and they enhance the current state of the art in RL research. It is possible to create intelligent, adaptable, and autonomous software agents that may improve human well-being by solving important problems and using the most recent developments in RL algorithms and techniques.

REFERENCES

[1]. Anon (2022). The Role of Reinforcement Learning in Autonomous Systems |. [online] www.interviewkickstart.com. Available at: https://www.interviewkickstart.com/blog/reinforcement-learning-autonomous-systems [Accessed 3 Mar. 2024].

[2]. Fadi AlMahamid (2022). Reinforcement Learning Algorithms: An Overview and Classification | IEEE Conference Publication | IEEE Xplore. [online] ieeexplore.ieee.org. Available at: https://ieeexplore.ieee.org/abstract/document/9569056 [Accessed 3 Mar. 2024].

[3]. Kiumarsi, B., Vamvoudakis, K.G., Modares, H. and Lewis, F.L. (2018). Optimal and Autonomous Control Using Reinforcement Learning: A Survey. IEEE Transactions on Neural Networks and Learning Systems, [online] 29(6), pp.2042–2062. doi:https://doi.org/10.1109/TNNLS.2017.2773458.

[4]. Padakandla, S. (2021). A Survey of Reinforcement Learning Algorithms for Dynamically Varying Environments. ACM Computing Surveys, 54(6), pp.1–25. doi:https://doi.org/10.1145/3459991.

[5]. Shah, V. (2020). Reinforcement Learning for Autonomous Software Agents: Recent Advances and Applications. Revista Espanola de Documentacion Cientifica, [online] 14(1), pp.56–71. Available at: https://redc.revistas-csic.com/index.php/Jorunal/article/view/155 [Accessed 3 Mar. 2024].

[6]. Singh, B., Kumar, R. and Singh, V.P. (2021). Reinforcement learning in robotic applications: a comprehensive survey. Artificial Intelligence Review, 07(08). doi:https://doi.org/10.1007/s10462-021-09997-9.

[7]. Zhang, T. and Mo, H. (2021). Reinforcement learning for robot research: A comprehensive review and open issues. International Journal of Advanced Robotic Systems, 18(3), p.172988142110073. doi:https://doi.org/10.1177/17298814211007305.