International Journal of Modern Achievement in Science, Engineering and Technology (IJSET)
ISSN 3023-459X
DOI: 10.63053/ijset.60

Deep Reinforcement Learning Multi-Agent Systems

Danial Karimzadeh

MA in Information Technology Management, Department of Information Technology Management, Faculty of Management, Islamic Azad University E-Campus, Tehran, Iran
PhD Student in Computer Engineering - Artificial Intelligence and Robotics, Department of Computer Engineering, Faculty of Engineering, Islamic Azad University Of Central Tehran Branch, Tehran, Iran

ARTICLE INFO

Keywords:
AI, MARL, DRL, Deep Reinforcement Learning, RL, Reinforcement Learning, Multi-Agent Systems

ABSTRACT

Deep Reinforcement Learning (DRL) is a subfield of artificial intelligence that combines reinforcement learning and deep neural networks to solve complex problems. In multi-agent systems, intelligent agents interact simultaneously within an environment, and their decisions affect each other's behavior. This paper examines Multi-Agent Reinforcement Learning (MARL), a significant branch of DRL, which is applied in systems with multiple agents having either common or conflicting goals. Key algorithms such as MADDPG, QMIX, and Mean Field RL, along with popular frameworks like TensorFlow, PyTorch, and Keras, are introduced. The applications of MARL in various domains, including economic systems, robotics, intelligent transportation, and resource management, are explored, and its advantages and disadvantages are discussed. Despite challenges such as high computational costs and limited scalability, MARL has the potential to drive significant innovations in technology and industry. The status of MARL in Iran and globally is analyzed, emphasizing the importance of investment and collaboration between academia and industry for advancement in this field. In conclusion, the paper highlights MARL's capability to solve complex problems and improve interactions, pointing to its potential in robotics, financial systems, and artificial intelligence

Deep reinforcement learning (DRL) is one of the subfields of artificial intelligence that solves complex problems by combining reinforcement learning and deep neural networks. In multi-agent systems, intelligent agents interact simultaneously in the environment and their decisions affect each other's behavior. (Liang et al., 2022).

This field has attracted the attention of many researchers and industries due to its capabilities in solving dynamic and complex problems. Deep reinforcement learning (DRL) is especially used to solve complex high-dimensional problems. One of the important subfields of DRL in recent years is multi-agent reinforcement learning (MARL), which is specifically used in systems with multiple agents with common or conflicting goals. Due to their complexity, dynamics, and multiple interactions, financial systems are a suitable environment for applying multi-agent reinforcement learning (MARL). In these systems, various agents (such as traders, financial institutions, and automated algorithms) compete or cooperate. MARL can help optimize investment strategies and risk management. (Zhang and Jin, 2019)

State the problem

In recent years, significant advances in the field of artificial intelligence, especially in the field of deep reinforcement learning (DRL), have brought about a major revolution in complex systems. In multi-agent systems, a group of intelligent agents interact in a common environment with similar or conflicting goals. These systems have many applications in fields such as robotics, strategic games, intelligent transportation and resource management. (Forster et al., 2018).

One of the basic challenges in multi-agent systems is to establish coordination and cooperation between agents. These agents must be able to make independent and intelligent decisions that are not only optimal but also do not conflict with the goals of other agents or environmental constraints. Also, learning in dynamic and uncertain environments with incomplete information is another obstacle to the development of deep reinforcement learning (DRL) algorithms for multi-agent systems. (Bosonio et al., 2008).

On the other hand, increasing the number of factors leads to an increase in computational complexity and a decrease in the stability of the learning process. Research shows that many existing algorithms in deep reinforcement learning (DRL) do not have the necessary scalability to handle large numbers of agents. This issue is especially important in real fields such as traffic management or control of energy distribution networks. (Lu et al., 2017).

In addition, one of the main challenges in this field is the conflict between the agents' goals. In environments where competition between agents is important, it is necessary to develop strategies to maintain a balance between cooperation and competition. This can lead to an increase in the complexity of learning models and the need to use effective communication protocols. (Silver et al., 2016).

Multi-agent reinforcement learning (MARL) is one of the important fields in artificial intelligence that deals with modelling multi-agent systems in complex and dynamic environments. This field combines machine learning, game theory, and optimal decision-making and has diverse applications in artificial intelligence, including strategic games, distributed systems, and complex simulations. In multi-agent environments, the decisions of each agent directly affect the outcomes of other agents. Therefore, designing models capable of managing these complex interactions is of great importance. It is still a big challenge. The existence of multiple agents can cause constant changes in the environment, which makes sustainable learning difficult. Agents need to communicate effectively to increase their efficiency. Have, but designing optimal communication protocols for scalable systems can bring many challenges. In many cases, agents may have conflicting goals (such as cooperation and competition at the same time), which require careful balancing. (Liang et al., 2022).

MARL can help improve the decision-making process in complex and multi-agent environments. Many real-world systems, including intelligent transportation, energy distribution networks, and strategic games, can be modelled as multi-agent environments. MARL can help create systems that are capable of interacting with humans or other machines. Autonomous systems, such as robots and driverless cars, can exploit MARL to collaborate in dynamic environments. (Zhang and Jin, 2019)

In Iran, the use of DRL in multi-agent systems is still in its early stages. Despite efforts in areas such as urban traffic management and smart energy grids, limitations in access to advanced hardware and highquality data have slowed progress in this area. However, there are significant opportunities to use this technology in various industries. On the other hand, advanced countries such as America and China have made significant achievements and developed advanced algorithms for practical applications with large investments in this field. (Open E, 2023)

Multi-agent reinforcement learning (MARL) plays an important role in the development of multi-agent technologies. In today's world, multi-agent systems such as self-driving cars, energy distribution networks, and crisis management require efficient MARL algorithms. Issues such as traffic management, communication network optimization and economic simulations require models that can manage cooperation and competition simultaneously. MARL can help increase stability and efficiency and be effective in optimizing resources, reducing costs, and improving the performance of multi-agent systems. However, the gap between research and practical applications still exists, as many MARL researches remain in the simulation stage and need to be developed for use in real environments. The use of deep reinforcement learning (DRL) in robotics is one of the most advanced methods for solving complex problems that require dynamic decision-making, coordination and interaction between robots. (Vinyl, 2019)

Multi-agent reinforcement learning (MARL) is used in the field of robotics to perform cooperation in complex missions, coordinate under uncertain conditions, and solve scalability issues. These algorithms are usually trained in simulated environments such as OpenAI Gym. However, in transferring to real environments, they face challenges such as noise, dynamic changes, and hardware limitations. (Yang et al., 2021)

Robots must be able to perform their tasks without the need for human intervention and with effective cooperation. However, the interdependence between robots and the lack of accurate communication can lead to a decrease in performance. In multi-robot systems, optimizing the allocation of resources such as batteries, tools, and workspace remains a major challenge. Existing algorithms are not designed for large numbers of robots and may be less efficient in larger environments. Robots can perform coordinated and effective missions such as search and rescue, infrastructure repairs, and environmental monitoring. MARL allows robots to be more autonomous in dangerous or inaccessible environments. (Zhang and Jin, 2019)

Robots can cooperate in carrying out tasks such as transporting objects, construction and moving loads in industrial environments. In times of crisis, a group of robots can help identify victims, send aid, and open paths by dividing tasks. In smart warehouses, robots can work together to move goods optimally. Also, the use of multi-agent reinforcement learning (MARL) is very effective in managing agricultural robots to perform tasks such as irrigation, planting, and harvesting. Financial systems are highly affected by unpredictable events, and modelling and predicting market behaviour is difficult due to its complexity and high noise. It is still challenging. (Yang et al., 2021)

In these environments, the strategy of each agent is highly dependent on the strategy of other agents, and MARL algorithms still face difficulties in managing these complex interactions. Real financial data is usually limited for security and economic reasons, which means that models can only be trained in simulated environments, and these simulations may not accurately reflect market behaviour. In financial systems, agents usually have conflicting goals (such as profit maximization versus risk mitigation) and balancing these goals remains a major challenge. (Vinyl, 2019)

MARL can help design automated trading algorithms that are capable of making decisions in complex situations. By using MARL models, it is possible to identify high-risk patterns in transactions and prevent financial crises. Also, MARL can be used to simulate market behaviour with multiple traders and test new economic policies. Agents can work together to design an optimal portfolio that will increase profits and reduce risk. (Liang et al., 2022)

Therefore, the need for more research on the development of DRL algorithms for multi-agent systems is felt so that they can solve the existing challenges and create wider practical applications at the national and international levels.

Challenges in multifactorial reinforcement learning

1. Lack of scalable algorithms that can handle a large number of agents.

2. Absence of comprehensive frameworks for the sustainability of learning in multi-agent systems.
3. The challenge in creating coordination between agents in conditions with high noise and uncertainty.
4. Lack of widespread applications in the real world due to computational limitations.
5. The challenge of complex interactions: how to optimize interactions between agents?
6. Environmental uncertainty: dynamic and unpredictable environments require flexible learning.
7. Conflicting goals of agents: balance between cooperation and competition among agents. (Zhang and Jin, 2019)
8. Need for coordination: development of efficient communication protocols between agents to achieve common goals.
9. Scalability issues: As the number of agents increases, the state and action space increase exponentially, which makes learning difficult. Existing algorithms are usually designed for a limited number of agents and are not efficient for large systems (such as smart cities or multi-robot networks).
10. Interdependence between agents: the behaviour of each agent depends on the strategies of other agents, which causes instability in learning.
11. Absence of common rewards: agents may have conflicting goals and each must consider their own optimization. Many real-world environments have multi-criteria or time-dependent rewards that DRL cannot handle well. (Liang et al., 2022).
12. Stable equilibrium problem: One of the basic challenges in MARL is to find a stable equilibrium (such as Nash equilibrium), especially in complex and nonlinear environments. Many existing algorithms cannot guarantee that the agents reach an equilibrium strategy.

.13Inconsistency in agent behaviour modelling: In real-world environments, agents often have irrational, random, or unknown behaviours that are difficult to model. (Vinyl, 2019)

Objectives in multifactorial reinforcement learning

Improving coordination and cooperation between agents in complex environments.

Simulation and use of MARL in real environments such as smart cities and energy management. (Zhang and Jin, 2019)

Advanced technology: DRL is recognized as a tool for enhancing artificial intelligence and automating systems.

Global Applications: From resource management to strategic games, DRL has a special place in various industries.

National development: Countries that invest in this area can gain a global competitive advantage.

Providing scalable algorithms for multi-agent systems and developing DRL algorithms suitable for complex and multi-agent environments (Liang et al., 2022)

Improving coordination and cooperation between agents in dynamic systems.

Increasing the efficiency of learning systems by reducing computational costs.

Creating real applications in fields such as intelligent transportation and robotics. (Vinyl, 2019)

Advanced MARL approaches

1. Methods based on cooperation:

These approaches are used for problems where agents must cooperate to achieve a common goal.

Shared Q-Learning: All agents share a common Q-Table.

Centralized Training and Decentralized Execution (CTDE):

Learning is done centrally, but each agent acts independently.

2. Competitive methods:

For environments where agents have competition, algorithms such as Nash-Q Learning and Self-Play are used. Using MARL to simulate competition between traders. Designing algorithms that can predict and adapt the behaviour of large traders (such as investment funds). (Zhang and Jin, 2019)

3. Blended learning:

These approaches are a combination of cooperation and competition, such as economic environments or war simulations.

4 . Central coordination:

In this approach, all robots are controlled by a central controller. They have high precision in coordination need strong communication and consume a lot of resources

5. Decentralized learning:

Each robot makes decisions independently. Greater flexibility and reduced need for communication are the challenges in coordination between robots.

6. Hybrid approach: a combination of centralized training and decentralized implementation. This approach is supported by algorithms such as MADDPG and QMIX. (vinyl, 2019)

Key algorithms in MARL

1. MADDPG (Multi-Agent Deep Deterministic Policy Gradient):

This algorithm is designed based on CTDE and uses Actor-Critic networks to improve stability.

2. QMIX:

In this method, a central neural network combines rewards and optimizes decision-making. (Zhang and Jin, 2019)

3. Mean Field RL

To reduce complexity in environments with a large number of agents, it models the average behaviour of other agents.

4. OpenAI Gym

To simulate learning environments.

5 .PyMARL

To test different MARL algorithms

6 . Ray Rllib

A scalable framework for implementing DRL and MARL algorithms. (Liang et al., 2022)

Multi-agent reinforcement learning frameworks

Deep learning frameworks are powerful tools that allow data researchers and engineers to quickly and easily build their deep learning models. Among the popular frameworks, the following can be mentioned:

TensorFlow

Is an open source library developed by Google and used for a wide range of machine-learning applications. TensorFlow has a large community and strong support and is well-suited for implementing complex deep-reinforcement learning models.

PyTorch

is another open-source library developed by Facebook. PyTorch has become very popular among researchers due to its simple and flexible user interface. PyTorch is great for research and rapid prototyping development.

Keras

is a high-level API for TensorFlow and Theano that allows users to quickly and easily build deep learning models. Keras is great for people new to deep learning. (Zhang and Jin, 2019)

Choosing the right algorithm and framework depends on various factors such as:

(1Environment complexity: For simpler environments, simpler algorithms such as DQN are sufficient. But for more complex environments, more advanced algorithms like PPO are needed.

(2Amount of data: If a large amount of training data is available, more complex models such as deep neural networks with many layers can be used.

(3Calculation time: Some algorithms, such as A3C, have a higher learning speed due to parallelization.

(4Computing equipment: Different frameworks may require different hardware. (Yang et al., 2021(

Applications in multifactorial reinforcement learning

Economic systems: modelling markets and simulating decision-making in financial environments.

Robotics: the use of multi-agent robots in search and rescue. Using MARL to coordinate multiple robots to perform complex tasks such as moving objects or navigating.

Smart transportation: optimizing the routes of self-driving cars.

Strategic games: developing intelligent agents for competitive games.

Resource management: optimal distribution of energy and natural resources. (Zhang and Jin, 2019)

Management of communication networks: optimization of data flow and resource allocation in wireless networks.

Multi-Agent Games: Used in games like 2Dota and StarCraft II to create intelligent agents

Bank liquidity management: optimizing the allocation of resources between banks to reduce the risk of financial crisis.

Algorithmic trading: MARL can design strategies that can optimize trades in real-time.

Risk management: Identifying high-risk scenarios in financial markets. Optimizing investments to reduce possible losses in critical situations. (Liang et al., 2022)

Market simulation: creating simulated environments in which the behaviour of traders and the effect of various economic policies are analyzed.

Portfolio management: MARL can help manage assets so that portfolios are tailored to meet financial goals and risk levels. (Vinyl, 2019)

Advantages of multifactorial reinforcement learning

High flexibility: quick adaptation to changes in the environment.

Ability to solve complex problems: managing dynamic and multidimensional environments.

Optimizing decision-making: improving efficiency and reducing costs. (Yang et al., 2021)

Disadvantages of multifactorial reinforcement learning

1. High computational cost: the need for advanced hardware such as GPU.
2. Limited scalability: the challenge in managing very large systems.
3. Dependence on simulated data: lack of real data for training. (Zhang and Jin, 2019)

The situation of multifactorial reinforcement learning in Iran

Challenges: Lack of research funding, limited access to advanced hardware, and low focus in this area.

Opportunities: DRL can be used in urban traffic management, energy networks and industrial systems. (Vinyl, 2019)

Global status of multifactorial reinforcement learning

Advances: Companies like DeepMind and OpenAI have come up with advanced algorithms.

Applications: self-driving cars, smart electric grids, and strategic games. (Yang et al., 2021)

Conclusion

Deep reinforcement learning in multi-agent systems can solve complex problems and optimize interactions. Despite challenges such as scalability and computational costs, appropriate investment and collaboration between academia and industry can yield significant results. Iran can also progress in this field by developing infrastructure and focusing on local applications. Multi-agent reinforcement learning is one of the prominent and advanced topics of DRL, which plays an important role in solving complex real-world problems. With further advancements in algorithms and tools, this field has the potential to create huge innovations in technology and industry. MARL in robotics has great potential to solve real-world challenges. With the development of scalable and noise-resistant algorithms, it is possible to achieve intelligent, coordinated and independent multi-robot systems that revolutionize various industries. MARL in financial systems enables the development of advanced models that can facilitate complex decision-making, simulate markets and improve financial performance. However, for the practical application of these models in the real world, further research is needed to overcome theoretical and practical challenges. MARL in artificial intelligence has the potential to create complex and autonomous systems that can operate in high-performance multi-agent environments. With the progress of research in this field, artificial intelligence systems will be able to solve more challenging and diverse problems.

References

1. Busoniu, L., Babuska, R., & De Schutter, B. (2008). A comprehensive survey of multi-agent reinforcement learning. IEEE Transactions on Systems, Man, and Cybernetics.
2. Foerster, J., Nardelli, N., Farquhar, G., & Whiteson, S. (2018). Counterfactual multi-agent policy gradients. Proceedings of the AAAI Conference on Artificial Intelligence.
3. Lowe, R., Wu, Y., Tamar, A., & Abbeel, P. (2017). Multi-agent actor-critic for mixed cooperative-competitive environments. Advances in Neural Information Processing Systems.
4. OpenAI. (2019). Solving Rubik's Cube with a robot hand.
5. OpenAI Research Blog. (2023). Latest advancements in deep reinforcement learning.
6. Silver, D. (2015). Deep reinforcement learning: Mastering the game of Go. Nature.
7. Silver, D., Schrittwieser, J., Simonyan, K., & Hassabis, D. (2019). Mastering the game of Go with deep neural networks and tree search. Nature.
8. Sutton, R. S., & Barto, A. G. (2018). Reinforcement learning: An introduction. MIT Press.
9. Vinyals, O., et al. (2019). Grandmaster level in StarCraft II using multi-agent reinforcement learning. Nature.
10. Yang, F., et al. (2021). Multi-agent reinforcement learning for AI applications