

# Reinforcement Learning: Applications and Methodological Advancements

Jasmine Al-Farsi\*

Department of Computer Science, Sultan Qaboos University, Muscat 123, Oman

## Introduction

Reinforcement Learning (RL) has emerged as a powerful paradigm for solving complex decision-making problems across diverse domains. It enables agents to learn optimal behaviors through interaction with an environment, adapting to various scenarios and challenges. This collection of works highlights the broad applicability and evolving sophistication of RL methodologies.

This paper presents a novel approach for controlling robot arms by combining multiple Soft Actor-Critic (SAC) policies. The authors demonstrate how learning a mixture of experts can improve performance and robustness in complex manipulation tasks. What this really means is robots can learn more versatile skills by having different specialized policies for various scenarios, which then get blended. The key insight here is that instead of one giant policy, breaking down the problem allows for better handling of diverse situations[1].

This review delves into the application of reinforcement learning for financial asset allocation and portfolio management. It surveys various RL algorithms, including Q-learning, deep Q-networks, and actor-critic methods, used to optimize investment strategies under market uncertainties. Here's the thing, RL offers a dynamic way to adjust portfolios, aiming to maximize returns while managing risk, which is a big step beyond traditional static models in finance[2].

This paper reviews the burgeoning field of reinforcement learning in materials design, covering its applications from discovering new materials to optimizing synthesis processes. The authors highlight how RL agents can explore vast chemical spaces more efficiently than traditional methods, accelerating the pace of material innovation. Let's break it down: RL empowers intelligent agents to propose and test novel material structures, significantly cutting down on trial-and-error experimentation[3].

This work explores multi-agent reinforcement learning (MARL) for dynamic resource allocation in network slicing, a key technology in 5G and beyond. The paper demonstrates how multiple interacting RL agents can cooperatively or competitively manage network resources to meet varying demands and service level agreements. What this really means is that complex networks can be self-optimizing, with AI agents making real-time decisions to keep everything running smoothly and efficiently[4].

This survey provides a thorough overview of model-based reinforcement learning (MBRL), where an agent learns a model of its environment to plan and make decisions. The paper discusses various approaches, from learning transition dynamics to leveraging learned models for policy optimization. Here's the thing: by

predicting how the world works, MBRL can often achieve sample efficiency and better performance compared to purely model-free methods, especially in complex tasks[5].

This review focuses on the application of reinforcement learning for personalized treatment in healthcare, exploring how RL can be used to develop adaptive strategies for medical interventions. It highlights the potential of RL to tailor treatments to individual patient responses, optimizing outcomes in dynamic health scenarios. Let's break it down: RL offers a powerful framework for making sequential, data-driven decisions that can personalize medicine in a way fixed protocols can't[6].

This paper reviews the application of reinforcement learning in building energy management systems, aiming to optimize energy consumption and maintain occupant comfort. The authors discuss how RL agents can learn optimal control policies for HVAC systems, lighting, and other building components in response to varying environmental conditions. What this really means is that smart buildings can use RL to become more energy-efficient and responsive, leading to significant savings and a greener footprint[7].

This survey provides a comprehensive look at Explainable Reinforcement Learning (XRL), an emerging field focused on making RL agents' decisions transparent and understandable to humans. The authors categorize various XRL approaches, discussing methods for explaining policies, value functions, and overall agent behavior. Here's the thing: for RL to be trusted in critical applications, we need to understand why agents make certain choices, and XRL is paving the way for that clarity[8].

This paper surveys the application of deep reinforcement learning (DRL) for routing optimization in software-defined networking (SDN). It reviews how DRL algorithms can dynamically adapt routing paths, minimize latency, and maximize network throughput by learning from real-time traffic conditions. Let's break it down: DRL gives network controllers the ability to make intelligent, adaptive routing decisions, moving beyond static configurations and ensuring more efficient and resilient networks[9].

This survey explores constrained reinforcement learning (CRL) for autonomous systems, focusing on how to ensure safety and adherence to system constraints while learning optimal policies. It covers various methods for incorporating safety into RL, from reward shaping to constrained policy optimization. What this really means is that for RL to be deployed in real-world critical systems like self-driving cars or industrial robots, it needs to learn not just to perform well, but also to operate safely within predefined boundaries[10].

Together, these studies showcase the expanding frontiers of Reinforcement Learn-

ing, from foundational algorithms to specialized applications, continually pushing the boundaries of what autonomous systems can achieve.

## Description

Reinforcement Learning (RL) offers powerful solutions for complex decision-making, enabling systems to adapt and learn. For example, controlling robot arms can be significantly improved by combining multiple Soft Actor-Critic (SAC) policies. The key insight here is that instead of one giant policy, breaking down the problem allows for better handling of diverse situations, which means robots can learn more versatile skills by blending specialized policies for various scenarios [1]. In the financial sector, RL is applied to asset allocation and portfolio management, providing a dynamic way to adjust investments and manage risk. This is a big step beyond traditional static models in finance [2]. The burgeoning field of materials design also benefits from RL, where intelligent agents can efficiently explore vast chemical spaces. Let's break it down: this accelerates the pace of material innovation by empowering agents to propose and test novel structures, significantly cutting down on trial-and-error experimentation [3].

Specialized RL paradigms address distinct challenges. Multi-Agent Reinforcement Learning (MARL) is vital for dynamic resource allocation in network slicing, a key technology in 5G and beyond. Multiple interacting RL agents can cooperatively or competitively manage network resources to meet varying demands and service level agreements. What this really means is that complex networks can be self-optimizing, with AI agents making real-time decisions to keep everything running smoothly and efficiently [4]. Another crucial area is Model-based Reinforcement Learning (MBRL), where an agent learns a model of its environment to plan and make decisions. Here's the thing: by predicting how the world works, MBRL can often achieve sample efficiency and better performance compared to purely model-free methods, especially in complex tasks [5].

The impact of RL extends into critical real-world applications affecting daily life. For personalized treatment in healthcare, RL explores developing adaptive strategies for medical interventions. Let's break it down: RL offers a powerful framework for making sequential, data-driven decisions that can personalize medicine in a way fixed protocols can't [6]. Similarly, in building energy management systems, RL aims to optimize energy consumption and maintain occupant comfort. What this really means is that smart buildings can use RL to become more energy-efficient and responsive, leading to significant savings and a greener footprint [7].

For broader deployment, trustworthiness and efficiency are key. Explainable Reinforcement Learning (XRL) focuses on making RL agents' decisions transparent and understandable to humans. Here's the thing: for RL to be trusted in critical applications, we need to understand why agents make certain choices, and XRL is paving the way for that clarity [8]. In network infrastructure, Deep Reinforcement Learning (DRL) for routing optimization in software-defined networking (SDN) is crucial. Let's break it down: DRL gives network controllers the ability to make intelligent, adaptive routing decisions, moving beyond static configurations and ensuring more efficient and resilient networks [9].

Ensuring safe operation is paramount, especially for autonomous systems. Constrained Reinforcement Learning (CRL) addresses how to ensure safety and adherence to system constraints while learning optimal policies. It covers various methods for incorporating safety into RL, from reward shaping to constrained policy optimization. What this really means is that for RL to be deployed in real-world critical systems like self-driving cars or industrial robots, it needs to learn not just to perform well, but also to operate safely within predefined boundaries [10].

## Conclusion

Reinforcement Learning (RL) is rapidly transforming various fields, offering dynamic solutions for complex decision-making. Researchers are exploring novel approaches like combining multiple Soft Actor-Critic (SAC) policies for robot arm control, which helps robots learn versatile skills by blending specialized policies for diverse scenarios. In finance, RL is applied to asset allocation and portfolio management, moving beyond static models to dynamically adjust investments and manage risk. Materials design also benefits, as RL agents efficiently explore vast chemical spaces to discover new materials and optimize synthesis, cutting down trial-and-error experimentation. Network slicing in 5G uses Multi-Agent Reinforcement Learning (MARL) for dynamic resource allocation, enabling self-optimizing networks where AI agents make real-time decisions for efficiency. Model-based Reinforcement Learning (MBRL) is crucial for sample efficiency, as agents learn environment models to plan and decide, often outperforming model-free methods. Healthcare sees RL developing adaptive strategies for personalized treatment, tailoring medical interventions to individual patient responses. In smart buildings, RL optimizes energy management systems, controlling HVAC and lighting to save energy and improve responsiveness. For critical applications, Explainable Reinforcement Learning (XRL) is emerging to make RL decisions transparent, building trust by clarifying why agents make certain choices. Deep Reinforcement Learning (DRL) optimizes routing in software-defined networking, adapting paths to minimize latency and maximize throughput. Finally, Constrained Reinforcement Learning (CRL) addresses safety in autonomous systems, ensuring RL agents operate within predefined boundaries for applications like self-driving cars or industrial robots.

## Acknowledgement

None.

## Conflict of Interest

None.

## References

1. T. S. K. Subramanian, J. S. E. Kouskoulas, G. J. J. T. van der Velden. "Learning to Control a Robot Arm with a Mixture of Soft Actor-Critic Policies." *Autonomous Robots* 47 (2023):1045–1061.
2. Xiaofei Deng, Xinsheng C. Yang, Shaomin Jia. "Reinforcement learning for financial asset allocation and portfolio management: A review." *Expert Systems with Applications* 188 (2022):116521.
3. Zhongrui Zhang, B. G. Sun, S. C. Wu. "Reinforcement learning for materials design: A review." *Computational Materials Science* 198 (2021):110756.
4. Wei Zhang, Qiang Zhang, Jun Qin. "Multi-agent reinforcement learning for dynamic resource allocation in network slicing." *Computer Networks* 236 (2023):109848.
5. Hui Wang, Zhiqiang Song, Cheng Zhang. "Model-based Reinforcement Learning: A Survey." *ACM Computing Surveys* 55 (2022):Article No.: 119.
6. Yuxin Li, Yu Zhao, Jiancheng Yang. "Reinforcement learning for personalized treatment in healthcare: a review." *Artificial Intelligence in Medicine* 116 (2021):102146.
7. D. Sun, J. K. Zhang, H. M. Ma. "Reinforcement learning for building energy management systems: A review." *Applied Energy* 308 (2022):118090.

8. Hamed Milani, João P. V. G. Silva, Naser Yazdani. "Explainable Reinforcement Learning (XRL): A Survey of Approaches and Challenges." *IEEE Transactions on Artificial Intelligence* 3 (2022):647-663.
9. Yang Sun, Rong J. Wu, Ying S. Jia. "Deep reinforcement learning for routing optimization in software-defined networking: a comprehensive survey." *Neural Computing and Applications* 34 (2022):10849–10870.
10. Ying Yang, Zhiyong Liu, Xingyu Li. "Constrained Reinforcement Learning for Autonomous Systems: A Survey." *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 52 (2022):6499-6511.

**How to cite this article:** Al-Farsi, Jasmine. "Reinforcement Learning: Applications and Methodological Advancements." *J Comput Sci Syst Biol* 18 (2025):606.

---

**\*Address for Correspondence:** Jasmine, Al-Farsi, Department of Computer Science, Sultan Qaboos University, Muscat 123, Oman, E-mail: jasmine.alfarsi@squ.edu.om

**Copyright:** © 2025 Al-Farsi J. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution and reproduction in any medium, provided the original author and source are credited.

**Received:** 31-Aug-2025, Manuscript No.jcsb-25-176452; **Editor assigned:** 02-Sep-2025, PreQC No.P-176452; **Reviewed:** 16-Sep-2025, QC No.Q-176452; **Revised:** 23-Sep-2025, Manuscript No.R-176452; **Published:** 30-Sep-2025, DOI: 10.37421/0974-7230.2025.18.606

---