RADemics

Multi-Objective Reinforcement Learning for Balancing Efficiency and Safety in Autonomous Decision Making

Afreen Kubra, Mekala
VISVESVARAYA TECHNOLOGICAL UNIVERSITY, VELALAR COLLEGE OF ENGINEERING AND TECHNOLOGY

13. Multi-Objective Reinforcement Learning for Balancing Efficiency and Safety in Autonomous Decision Making

[1]Afreen Kubra, Assistant Professor, Department of Artificial Intelligence and Machine Learning, HKBK College of Engineering, Visvesvaraya Technological University, Bangalore 560045, Karnataka, India, afteenkubra96@gmail.com

[2]S. Mekala, Assistant Professor, Department of Information Technology, Velalar College of Engineering and Technology, Erode, Tamil Nadu, India, mekala2712@gmail.com

Abstract

This book chapter explores the integration of multi-objective reinforcement learning (MORL) for balancing efficiency and safety in autonomous decision-making. As autonomous systems are increasingly deployed in complex, real-world environments, ensuring both high-performance optimization and safety becomes paramount. The chapter delves into the fundamental concepts of MORL, emphasizing how multiple conflicting objectives can be addressed simultaneously, while maintaining robust safety standards. Various methods, such as risk-sensitive learning, safety constraints, and reward shaping, are examined to effectively guide autonomous agents in navigating trade-offs between exploration and safety. Additionally, the chapter discusses policy-based methods, verification, and validation techniques crucial for real-world implementation. With a focus on real-world testing for safety assurance, it presents strategies for integrating safety-critical components in decision-making processes. This work provides valuable insights into enhancing the reliability, performance, and safety of autonomous systems in dynamic environments.

Keywords: Multi-Objective Reinforcement Learning, Safety Assurance, Autonomous Decision-Making, Reward Shaping, Risk-Sensitive Learning, Policy Convergence.

Intrtoduction

Autonomous systems have become central to a variety of industries, from self-driving cars and drones to robotics and industrial automation [1,2]. The primary challenge in deploying these systems lies in balancing two critical objectives: optimizing for performance (efficiency) and ensuring safety [3,4]. Multi-objective reinforcement learning (MORL) offers a promising framework to address this challenge, as it can simultaneously optimize multiple objectives with potentially conflicting goals [5,6]. In the context of autonomous decision-making, safety often cannot be sacrificed for efficiency, making it essential to incorporate safety constraints into learning algorithms [7,8]. As these systems operate in dynamic and unpredictable environments, ensuring their robustness while meeting operational goals was essential for their widespread adoption [9,10].

One of the key advantages of MORL lies in its ability to model and balance multiple objectives, such as speed, cost, and safety, which are often at odds in real-world applications [11,12]. In traditional reinforcement learning (RL), a single objective was typically optimized, whichlead to risky behavior if safety was not explicitly integrated [13]. MORL, on the other hand, allows agents to consider multiple factors during decision-making, adjusting their strategies based on the relative importance of each objective [14,15]. This was particularly relevant for autonomous systems, where a failure to consider safety could result in catastrophic outcomes [16]. By exploring MORL techniques, this chapter aims to provide insights into the methods that can balance performance and safety in real-world autonomous systems [17].

Ensuring safety in autonomous decision-making was a complex task that extends beyond the traditional exploration-exploitation dilemma found in typical reinforcement learning [18]. While exploration was necessary for learning new information, it introduces risks that could compromise safety [19,20]. Autonomous agents must, therefore, explore and learn in a manner that minimizes potential harm [21]. To this end, various safety-preserving techniques, such as risk-sensitive learning, safety constraints, and reward shaping, are often incorporated into MORL models [22]. Risk-sensitive learning adjusts the decision-making process by considering not only the expected rewards but also the potential risks associated with each action [23]. By quantifying risks, autonomous agents can make more informed decisions that prioritize safety while still striving for efficiency [24].

A major challenge in multi-objective reinforcement learning was the design of reward functions that can effectively capture the trade-offs between different objectives [25]. Reward shaping, which modifies the reward structure to guide the agent's behavior towards desired outcomes, plays a vital role in achieving a balance between efficiency and safety. In safety-critical applications, such as autonomous driving or medical robotics, the agent’s reward function needs to strongly emphasize avoiding dangerous states. On the other hand, the agent must also be incentivized to explore and learn in order to optimize efficiency. The development of multi-objective reward functions that balance these needs was a critical area of research in MORL, as it directly influences the effectiveness of the learning process.