RADemics

MARKOV DECISION PROCESSES FOR MODELING SEQUENTIAL DECISION MAKING IN ROBOTICS

Rajan Singh, Nidhi Tiwari, B. Arun
IFTM UNIVERSITY, IFTM UNIVERSITY, RATHINAM TECHNICAL CAMPUS

MARKOV DECISION PROCESSES FOR MODELING SEQUENTIAL DECISION MAKING IN ROBOTICS

[1]Rajan Singh, Associate Professor, Department of Mathematics, School of Sciences, IFTM University, Moradabad, Uttar Pradesh – 244102. rajan0779@rediffmail.com

[2]Nidhi Tiwari, Assistant Professor, Department of Mathematics, School of Sciences, IFTM University, Moradabad, Uttar Pradesh – 244102. nidhitiwari26@gmail.com

[3]B. Arun, Associate Professor, Department of Science and Humanities (Mathematics), Rathinam Technical Campus, Eachanari, Coimbatore -641021. arun.naam@gmail.com

Abstract

Autonomous robotic systems increasingly operate in dynamic and uncertain environments where intelligent action selection depends on the ability to evaluate long-term consequences rather than isolated decisions. Sequential decision making has therefore emerged as a fundamental research area for enabling adaptive perception, planning, navigation, manipulation, and collaborative task execution across diverse robotic applications. Among the available mathematical frameworks, Markov Decision Processes (MDPs) provide a rigorous probabilistic model for representing state transitions, action selection, reward optimization, and policy generation under uncertainty. The mathematical principles underlying MDPs establish the foundation for dynamic programming, optimal control, and reinforcement learning, enabling robotic systems to identify decision policies that maximize cumulative long-term rewards while adapting to continuously evolving operational conditions. This chapter presents a comprehensive examination of MDPs as a unified framework for modeling sequential decision making in robotics by integrating theoretical concepts with computational methodologies and practical implementation strategies. Fundamental principles of sequential decision making, mathematical formulations of MDPs, Bellman equations, value functions, policy representation, transition dynamics, and reward structures are systematically discussed to establish the theoretical basis of probabilistic robotic planning. Classical solution algorithms, including value iteration, policy iteration, approximate dynamic programming, and computational complexity analysis, are examined to highlight their convergence characteristics and scalability in large decision spaces. The chapter further explores MDP-based modeling for mobile robot navigation, robotic manipulation, autonomous exploration, and multi-robot coordination, demonstrating the versatility of probabilistic decision models across heterogeneous robotic platforms. The relationship between MDPs and reinforcement learning is analyzed to illustrate the evolution of intelligent robotic decision-making from model-based optimization toward data-driven autonomous learning. Emerging developments involving hierarchical decision models, safe reinforcement learning, explainable artificial intelligence, digital twins, edge intelligence, cloud robotics, and multi-agent autonomous systems are also reviewed to emphasize future research opportunities.

Keywords: Markov Decision Processes; Sequential Decision Making; Autonomous Robotics; Reinforcement Learning; Dynamic Programming; Multi-Robot Coordination

Introduction

The rapid evolution of robotics has transformed autonomous systems from executing predefined mechanical operations into intelligent platforms capable of making adaptive decisions in dynamic and uncertain environments [1]. Modern robotic systems are increasingly deployed in industrial automation, healthcare, logistics, agriculture, autonomous transportation, space exploration, disaster response, and service applications, where continuous interaction with complex surroundings demands advanced computational intelligence [2]. Every robotic operation involves perception, environmental interpretation, planning, action selection, and feedback, requiring the integration of sensing technologies with intelligent decision-making algorithms [3]. Environmental uncertainty caused by sensor noise, changing operating conditions, communication delays, and unpredictable external disturbances significantly complicates autonomous decision making. Conventional rule-based control methods often encounter limitations when addressing continuously evolving environments because predetermined responses cannot adequately represent the diversity of real-world scenarios. Mathematical models capable of reasoning under uncertainty have therefore become essential for improving robotic autonomy [4]. Such developments have established intelligent decision-making frameworks as a fundamental component of modern robotics, enabling autonomous systems to operate efficiently while adapting to complex operational conditions through systematic planning and long-term optimization [5].

Sequential decision making has emerged as one of the most significant paradigms for designing autonomous robotic systems capable of operating over extended periods in uncertain environments [6]. Unlike isolated decision-making processes that focus exclusively on immediate actions, sequential decision making considers the long-term consequences associated with every selected action. Each decision modifies the current operating state and influences future planning opportunities, creating an interconnected sequence of actions that collectively determines mission performance [7]. Robotic systems performing autonomous navigation, object manipulation, environmental exploration, and cooperative task execution continuously evaluate alternative decisions while balancing efficiency, safety, energy consumption, and task completion objectives. This long-term perspective enables robots to optimize cumulative performance rather than maximizing immediate rewards [8]. As robotic applications become increasingly sophisticated, computational frameworks capable of modeling state evolution, uncertainty, and future outcomes have become indispensable [9]. Sequential decision-making principles therefore provide the conceptual foundation for autonomous planning, adaptive control, and intelligent behavior across a wide range of robotic applications [10].

Among the mathematical models developed for sequential decision making, Markov Decision Processes (MDPs) provide one of the most comprehensive and theoretically rigorous frameworks for representing stochastic decision problems [11]. An MDP models the interaction between an autonomous agent and its environment using states, actions, transition probabilities, reward functions, and decision policies. The Markov property simplifies sequential optimization by assuming that future state transitions depend only on the present state and the selected action, eliminating the need to retain complete historical information [12]. This formulation enables efficient representation of uncertain environments while preserving the essential information required for optimal planning [13]. Bellman equations further strengthen the mathematical foundation of MDPs by introducing recursive optimization methods that estimate long-term cumulative rewards associated with alternative decisions [14]. Dynamic programming algorithms derived from these principles provide systematic procedures for identifying optimal policies capable of maximizing expected performance across multiple decision stages. Such characteristics have established MDPs as a fundamental mathematical framework supporting intelligent robotic decision making [15].