RADemics

Integration of Deep Reinforcement Learning and Evolutionary Strategies for Robotic Path Planning

Sneha Prakash, Anjali Sankar
ADI SHANKARA INSTITUTE OF ENGINEERING AND TECHNOLOGY, ADI SHANKARA INSTITUTE OF ENGINEERING AND TECHNOLOGY

Integration of Deep Reinforcement Learning and Evolutionary Strategies for Robotic Path Planning

[1]sneha Prakash, Assistant professor, Computer applications, Adi shankara institute of engineering and technology, kalady Mail ID: sneha.cs@adishankara.ac.in

[2]Anjali Sankar, Assistant Professor, Department of Computer Applications, Adi Shankara Institute of Engineering and Technology. Vidya Bharathi Nagar, Mattoor, Kalady, Kerala 683574, Mail ID: anjal.cs@adishankara.ac.in

Abstract

Robotic path planning in complex and dynamic environments remains a significant challenge due to the need for real-time adaptability, environmental uncertainty, and computational efficiency. Traditional algorithms often fail to generalize across varied terrains and obstacle configurations, necessitating more intelligent and adaptive solutions. This chapter presents a hybrid framework that integrates Deep Reinforcement Learning (DRL) with Evolutionary Strategies (ES) to develop robust, scalable, and generalizable navigation policies for autonomous robotic agents. The proposed integration leverages the exploration capacity of ES and the fine-tuned learning efficiency of DRL to enhance both convergence speed and behavioral resilience in unpredictable scenarios. Key innovations include the use of modular policy architectures, transferable policy graphs, and meta-learning techniques that enable the system to rapidly adapt across multi-environment navigation tasks. Extensive evaluation in simulated and varied settings demonstrates superior performance in terms of path optimality, generalization, and real-time decision-making compared to conventional standalone approaches. The chapter also introduces a set of architecture-aware evaluation frameworks, including ablation studies and environment-driven policy adaptation metrics, to measure generalization capability comprehensively. By bridging model-free learning and evolutionary optimization, the hybrid paradigm outlined in this work establishes a new foundation for scalable autonomous navigation in diverse and uncertain environments.

Keywords: Robotic Path Planning, Deep Reinforcement Learning, Evolutionary Strategies, Hybrid Intelligence, Policy Generalization, Autonomous Navigation

Introduction

The ability of robots to autonomously navigate through complex, dynamic, and partially known environments is foundational to their successful deployment in real-world applications such as autonomous driving, search and rescue missions, planetary exploration, and warehouse logistics [1]. Robotic path planning, therefore, remains a critical area of research in autonomous systems, focused on generating efficient and collision-free trajectories from an origin to a destination while accounting for environmental constraints and uncertainties [2]. Classical algorithms such as A*, Dijkstra’s, and RRT variants have provided elegant and computationally tractable solutions in deterministic or semi-structured environments [3]. As robotic platforms evolve toward greater autonomy and are expected to function in diverse, unstructured, and dynamic real-world domains, these conventional methods often fall short in terms of adaptability, robustness, and generalization [4]. This growing complexity in operational contexts calls for learning-driven, intelligent path planning paradigms capable of real-time decision-making under uncertainty [5].

Deep Reinforcement Learning (DRL) has recently emerged as a powerful solution to the challenges faced in robotic navigation by enabling agents to learn optimal policies through interactions with their environment [6]. By combining the representational power of deep neural networks with reinforcement learning's sequential decision-making capabilities, DRL facilitates end-to-end learning from high-dimensional sensory inputs to control actions [7]. Its ability to learn directly from raw inputs such as images, lidar scans, or point clouds eliminates the need for manual feature engineering, making it particularly suited for high-complexity environments [8]. Nonetheless, DRL models often face significant challenges in real-world applications. These include poor sample efficiency, instability during training, sensitivity to hyperparameters, and a tendency to overfit to training environments [9]. The reliance on finely tuned reward functions and the difficulty in handling sparse rewards exacerbate the problem of generalization, making standalone DRL approaches inadequate for broad deployment across varied and unseen settings [10].

To overcome these limitations, Evolutionary Strategies (ES) have gained attention as an alternative or complementary approach. ES are population-based optimization techniques inspired by biological evolution, operating through selection, mutation, and recombination of candidate solutions across generations [11]. These strategies are inherently parallelizable, require no gradient information, and are well-suited for navigating noisy, non-differentiable, or deceptive search spaces [12]. In robotic path planning, ES can be used to evolve control policies, optimize hyperparameters, and explore diverse behavioral strategies [13]. Their robustness to environmental perturbations and their capacity to maintain diversity in the policy population make them particularly attractive for applications involving uncertainty and non-linearity [14]. ES alone can struggle with fine-grained policy learning and may require extensive computational resources to converge. This motivates the integration of DRL and ES into unified hybrid frameworks that combine the respective strengths of local policy refinement and global exploration [15].