

The logo consists of a blue arrow pointing to the right, containing the text "RADemics" in white. To the left of the arrow is a thick, dark blue vertical bar. At the bottom left, there are several thin, curved lines in shades of blue and grey, resembling stylized grass or reeds.

RADemics

Reinforcement Learning Algorithms for Control Optimization in Embedded and IoT Based Electrical Systems

Ayyasamy Thangaraj, S. Ohmshankar,
Varsha Dange

ACS COLLEGE OF ENGINEERING, AGNI COLLEGE
OF TECHNOLOGY, VISHWAKARMA INSTITUTE OF
TECHNOLOGY

Reinforcement Learning Algorithms for Control Optimization in Embedded and IoT Based Electrical Systems

¹Ayyasamy Thangaraj, Associate Professor, Aerospace Engineering, ACS College of Engineering, No.207, Kambipura, Mysore Road, Bangalore -74, ayyasamy@acsce.edu.in

²S. Ohmshankar, AP, ECE, Agni College of Technology, ohmshankar.ece@act.edu.in, ohmshankar.ece@gmail.com

³Varsha Dange, Assistant professor, CSE in AIML, Vishwakarma Institute of Technology, Pune-411037, India. varsha.dange1@vit.edu

Abstract

The integration of Reinforcement Learning (RL) into embedded systems and IoT-based electrical networks has sparked significant advancements in control optimization, offering potential solutions to enhance the efficiency and adaptability of modern systems. However, the application of RL in such resource-constrained environments faces critical challenges related to energy efficiency, computational constraints, and real-time performance. This chapter explores innovative hybrid RL methods and hardware-software co-design strategies to minimize power consumption while maintaining optimal system performance in embedded devices. Key techniques such as energy harvesting, power management, and adaptive control policies are presented, highlighting their role in overcoming the limitations of low-power IoT devices. Additionally, the chapter delves into the challenges of model compression, resource management, and latency in real-time systems, proposing solutions to ensure continuous, autonomous operation without compromising system functionality. The evolving landscape of RL in embedded and IoT systems is thoroughly examined, offering insights into future research directions that can drive the next generation of energy-efficient, intelligent systems for a wide range of applications, from smart grids to autonomous transportation.

Keywords: Reinforcement Learning, Energy Efficiency, Embedded Systems, IoT, Real-Time Performance, Power Management

Introduction

The exponential growth of Internet of Things (IoT) and embedded systems has revolutionized a wide range of industries, including healthcare, industrial automation, transportation, and smart cities [1]. These systems often operate under strict constraints, such as limited power supply, processing capabilities, and memory [2]. The need for intelligent, autonomous decision-making in real-time control systems has brought Reinforcement Learning (RL) to the forefront as a promising technique [3]. RL enables systems to learn optimal actions based on feedback from the environment, improving their performance over time without explicit programming [4]. Deploying RL in embedded systems and IoT networks comes with substantial challenges, especially when

energy efficiency and computational resources are limited. This chapter delves into the growing importance of RL in these contexts and explores various strategies to enhance the energy efficiency of RL-based systems while ensuring real-time performance [5].

A major hurdle in applying RL to embedded systems is the energy consumption required to process large amounts of data and make decisions in real-time [6]. Traditional RL models, particularly Deep Reinforcement Learning (DRL), require substantial computational resources and power, making them unsuitable for low-power devices [7]. For IoT applications, which often run on battery-powered sensors or energy-harvesting systems, power management is critical [8]. Energy consumption must be minimized without compromising the effectiveness of the learning algorithm. The challenge lies in adapting RL models to be resource-efficient while still performing complex tasks, such as controlling industrial machines, managing traffic flows in smart cities, or optimizing energy usage in smart grids [9]. This section explores these energy-efficiency challenges and emphasizes the need for new approaches that integrate power-aware learning algorithms into resource-constrained environments [10].

In response to these challenges, several hybrid approaches have emerged, combining traditional control methods with RL to achieve the best of both worlds [11]. Techniques like Model Predictive Control (MPC), PID control, and other classical control theories offer stability and predictable performance, while RL adds adaptability and optimization capabilities [12]. These hybrid approaches can be tailored to reduce the computational burden of RL algorithms by relying on traditional models for routine tasks and using RL to optimize more complex decision-making processes [13]. This hybridization helps minimize power consumption by reducing the frequency and intensity of RL model updates [14]. The integration of these two paradigms offers a balanced approach to meet the needs of real-time applications in energy-constrained environments. The synergy between classical and RL techniques allows for more efficient decision-making in embedded systems, where energy is a premium resource [15].