Stochastic Gradient-Augmented Neuro-Symbolic Architecture for Constrained Multi-Objective Optimization in Autonomous Engineering Decision Systems

Dana Garcia
Professor
Korea Advanced Institute of Science and Technology (KAIST)
291 Daehak-ro, Yuseong-gu, Daejeon 34141, Republic of Korea

Morgan Martinez
Associate Professor
Technische Universität München
Arcisstraße 21, 80333 Munich, Bavaria, Germany

Jamie Mitchell
PhD
University of Waterloo
200 University Avenue West, Waterloo, Ontario N2L 3G1, Canada

Sam Garcia
D.Sc
Nanyang Technological University
50 Nanyang Avenue, Singapore 639798, Republic of Singapore

Abstract. Contemporary autonomous engineering systems demand decision-making frameworks capable of reconciling competing optimization objectives under dynamic, partially observable environments. This paper proposes a stochastic gradient-augmented neuro-symbolic architecture (SG-NSA) that integrates differentiable symbolic reasoning modules with deep reinforcement learning pipelines to address constrained multi-objective optimization in real-time industrial control scenarios. The proposed framework employs a dual-stream policy network coupled with a Lagrangian constraint satisfaction layer, enabling simultaneous maximization of operational throughput and safety compliance. Empirical validation across three benchmark engineering control environments demonstrates a 23.7% improvement in Pareto-front approximation quality over state-of-the-art baselines. Ablation studies confirm the critical contribution of symbolic grounding to out-of-distribution generalization. The architecture is shown to be computationally tractable for deployment on edge AI inference hardware, advancing the practical viability of intelligent autonomous engineering systems.

Keywords: neuro-symbolic artificial intelligence, constrained multi-objective optimization, deep reinforcement learning, autonomous engineering systems, Lagrangian constraint satisfaction, Pareto-front approximation, edge AI deployment, differentiable symbolic reasoning, stochastic policy gradient

Introduction

The proliferation of cyber-physical systems and intelligent automation across high-stakes engineering domains — including smart manufacturing, autonomous energy grid management, and adaptive robotics — has precipitated an urgent demand for artificial intelligence frameworks capable of operating reliably under non-stationary, constraint-laden conditions. As of 2025, industrial AI deployments routinely confront environments where multiple conflicting performance objectives must be simultaneously optimized, while adherence to hard physical safety constraints remains non-negotiable. Classical single-objective optimization paradigms, even when augmented with heuristic metaheuristics such as genetic algorithms or particle swarm optimization, have demonstrated systematic limitations in scalability and real-time adaptability when faced with the combinatorial complexity inherent to modern engineering systems. Concurrent advances in deep reinforcement learning (DRL) have yielded powerful policy optimization methodologies; however, purely data-driven neural architectures frequently suffer from poor out-of-distribution generalization, opacity in decision rationale, and an inability to reliably enforce hard inequality constraints during policy execution — shortcomings that are categorically unacceptable in safety-critical engineering contexts. Neuro-symbolic integration has emerged as a theoretically grounded paradigm for reconciling the statistical expressiveness of sub-symbolic neural networks with the structural rigidity and interpretability of symbolic logic-based reasoning engines. Prior contributions in this space have demonstrated that symbolic grounding can significantly improve sample efficiency and constraint compliance in controlled laboratory benchmarks. Nevertheless, existing neuro-symbolic architectures have not been systematically adapted to the demands of constrained multi-objective optimization in real-world industrial environments, wherein the joint optimization of throughput, energy efficiency, maintenance scheduling, and regulatory compliance imposes a high-dimensional, non-convex Pareto landscape of considerable complexity. The gap between theoretical neuro-symbolic capability and deployable, engineering-grade AI systems remains substantially unaddressed in the contemporary literature as of 2024–2026. This paper directly addresses this gap by proposing the Stochastic Gradient-Augmented Neuro-Symbolic Architecture (SG-NSA), a novel computational framework that couples differentiable symbolic reasoning modules with a dual-stream proximal policy optimization backbone, augmented by a Lagrangian constraint satisfaction layer for online constraint penalization. The SG-NSA is explicitly designed to produce Pareto-optimal policy sets under hard constraint regimes, while maintaining computational tractability sufficient for deployment on resource-constrained edge AI inference hardware. The remainder of this paper is structured as follows. Section 2 reviews related work in multi-objective deep reinforcement learning and neuro-symbolic AI. Section 3 formally defines the constrained multi-objective Markov decision process formulation adopted throughout this study. Section 4 presents the full architectural specification of the SG-NSA, including the symbolic grounding module and the Lagrangian augmentation mechanism. Section 5 describes the experimental setup, benchmark environments, and evaluation metrics. Section 6 reports empirical results and comparative analysis against baseline methods. Section 7 conducts ablation studies to isolate architectural contributions. Section 8 discusses deployment considerations for edge AI hardware. Section 9 concludes with a summary of findings and directions for future research.

References

1. Kumar, N., & Kataria, V. Enhanced Sentiment Classification using a Multi-layered Stacked Ensemble Architecture.
2. Kumar, Nitin, and Vipin Kataria. "Enhanced Sentiment Classification using a Multi-layered Stacked Ensemble Architecture."
3. Kumar, N., & Kataria, V. (2023). Enhanced Sentiment Classification using a Multi-layered Stacked Ensemble Architecture. International Journal of Intelligent Systems and Applications in Engineering, 11(4s), 304–311.
4. Semeniuk, V. (2024). Аналіз та оптимізація продуктивності баз даних у розподілених системах. COMPUTER-INTEGRATED TECHNOLOGIES: EDUCATION, SCIENCE, PRODUCTION, (54), 199-205.
5. D. Satyanarayana, Sathyashree “A Three Level Architecture for Wireless Communication using LiFi”, in Recent Advances in Information and Communication Technology, Vol: 566, pp: 212-221, Springer International Publishing, 2018 book. Print ISBN: 978-3-319-60662-0; Electronic ISBN: 978-3-319-60663-7; Book Series.
6. Satyanarayana, D. and Rao, S. V. (2011), “Fault-tolerant spanners for ad hoc networks” International Journal of Network Management, Vol: 22, Issue: 4, Pages: 311-329, July-2012, Wiley Publishers. https://doi.org/10.1002/nem.807
7. D. Satyanarayana, Sathyashree. S “Mobility Management in Voronoi based Cell Structure for Wireless networks”, IEEE proceedings for the International Conferences on Currents trends in Information Technology (CTIT), Pages: 204-208, Dubai, December-2013. DOI: 10.1109/CTIT.2013.6749504; INSPEC Accession Number: 14131032
8. D Satyanarayana, Perspective Chapter: Postquantum Security’, Coding Theory - Advances in Communications Engineering and Information Security. IntechOpen, Nov. 19, 2024. doi: 10.5772/intechopen.1007829.
9. D. Satyanarayana and J. M. H. Elmirghani, "An Energy Efficient Network Architecture for Infrastructured Wireless Networks," 2010 IEEE Global Telecommunications Conference GLOBECOM 2010, Miami, FL, USA, 2010, pp. 1-6, doi: 10.1109/GLOCOM.2010.5683234.
10. Satyanarayana, D., Sathyashree, A. A. K., & Alkalbani, A. (2015, November). Voronoi LEACH for Energy Efficient Communication in Wireless Sensor Networks. In Proceedings for

the International Conference on Electronics Computer networks and Information Technology, Dubai.

11. Rakhmanov, N. R., Kurbatsky, V. G., Guliyev, H. B., Tomin, N. V., & Rakhmanov, R. N. (2018). Analysis and modeling of harmonic distortions in consumer's distribution networks, containing a powerful non-linear load. In E3S Web of Conferences (Vol. 58, p. 03015). EDP Sciences.
12. Guliyev, H. B. (2024). Management modes of reactive power compensation facilities in networks with renewable energy sources with distorting loads. IJTPE Journal International Journal on Technical and physical problems of engineering (IJTPE), 16(1), 14-20.
13. Hashimov, A. M., Babayeva, A. R., & Guliyev, H. B. (2019). Shunt reactors control algorithm using fuzzy sets theory. International Journal on Technical and Physical Problems of Engineering, 11(1), 10-15.