

Generative policy models for autonomous governance in edge AI

Awolesi Abolanle Ogunboyoy *

Independent researcher.

World Journal of Advanced Research and Reviews, 2025, 27(01), 1394-1398

Publication history: Received on 01 June 2025; revised on 08 July 2025; accepted on 10 July 2025

Article DOI: <https://doi.org/10.30574/wjarr.2025.27.1.2608>

Abstract

As AI increasingly migrates to edge environments characterized by decentralization, limited resources, and real-time demands, the need for autonomous governance mechanisms has become paramount. This study introduces Generative Policy Models (GPMs), a novel class of transformer-based generative reinforcement learning frameworks designed for self-evolving policy generation in edge AI settings. By synthesizing policies without reliance on labeled data or central supervision, GPMs enable autonomous swarms, adaptive IoT networks, and mission-critical edge systems to operate efficiently and intelligently. Furthermore, three simulated environments, UAV swarms, smart traffic control, and IoT resource allocation, were used to evaluate GPM performance. Results demonstrate that GPMs surpass traditional RL baselines in decision latency, adaptability, and policy novelty, confirming their suitability for real-world decentralized systems. This work fills a critical gap in the literature by merging generative AI with edge autonomy and paves the way for resilient, explainable, and self-governing AI infrastructures.

Keywords: Generative reinforcement learning; Edge AI; Autonomous governance; Transformer models; Decentralized policy; Self-evolving systems

1. Introduction

The evolution of Artificial Intelligence (AI) has gradually shifted from centralized, data-intensive environments to decentralized, resource-constrained edge computing infrastructures (Duan et al., 2022; Walia et al., 2023). In an era where edge AI governs autonomous drones, vehicular networks, and smart IoT environments, the central question arises: How can such distributed systems govern themselves autonomously without continuous cloud or human supervision?

As edge devices increasingly make autonomous decisions in volatile and mission-critical environments, self-evolving governance strategies become paramount (Kurt et al., 2022). This research introduces and evaluates a novel framework for Generative Policy Models (GPMs) that leverage generative reinforcement learning (GRL) to autonomously construct, adapt, and evolve decision policies at the edge. The demand for autonomous governance stems from the inability of static rule-based systems to adapt to the dynamic and unpredictable conditions typical of edge deployments (Golpayegani et al., 2024).

Moreover, traditional cloud-dependent governance strategies are infeasible, with data privacy, latency constraints, and connectivity instability inherent in edge architectures (Rajapakse et al., 2023). Generative policy models address these challenges by utilizing self-supervised reinforcement learning to generate optimal policies without relying on extensive labeled datasets or cloud intervention.

The following hypotheses drive this research:

* Corresponding author: Awolesi Abolanle Ogunboyoy

- Generative reinforcement learning models can autonomously learn governance policies in edge environments with minimal supervision,
- Such models are capable of evolving adaptively in response to dynamic operational contexts, and
- GPMs outperform traditional reinforcement learning and static policy baselines in edge-based autonomous decision-making.

By integrating policy generation with lightweight transformer architectures and edge-specific constraints such as computational load and latency thresholds, as researched by Shuvo et al. (2022) and Li et al. (2023), this work contributes a strategic leap in enabling autonomous swarms of devices and self-adaptive systems in the Internet of Things (IoT). Also, it investigates the feasibility, performance, and generalizability of GPMs within edge AI scenarios, thereby outlining their implications for self-regulating edge networks, smart city frameworks, and critical infrastructure monitoring systems.

2. Literature Review

The intersection of edge computing and AI governance has seen rapid growth, driven by the proliferation of intelligent devices and the need for low-latency decision-making (Bourechak et al., 2023; Hemmati et al., 2024). However, traditional edge AI models rely on supervised learning frameworks or cloud-supported policy management (Hussain et al., 2020); recent advances have highlighted the importance of decentralization and autonomy in policy generation (Cao, 2022; Zhu et al., 2024).

Edge environments, typified by limited connectivity, heterogeneous architectures, and high variability, necessitate adaptable and self-regulating policy frameworks. Generative models, particularly Variational Autoencoders (VAEs), Generative Adversarial Networks (GANs), and transformer-based architectures like GPT, have revolutionized many domains, including natural language processing, image synthesis, and protein structure prediction (Bengesi et al., 2024).

When fused with reinforcement learning (RL), these generative systems form a new paradigm known as Generative Reinforcement Learning (GRL), capable of synthesizing policies based on learned representations rather than predefined datasets (Cao et al., 2023; Chen et al., 2024). Furthermore, according to Chen et al. (2024), policy learning through deep reinforcement learning (DRL) has been explored in edge computing scenarios to optimize energy usage, task offloading, and network scheduling. However, these approaches often require frequent retraining and do not generalize well across devices or tasks. Meta-RL and federated RL have attempted to address these gaps by enabling knowledge transfer across agents, but the generative capacity for novel policy emergence remains under-explored (Hu et al., 2023; Wu et al., 2024).

This study identifies a key gap in the existing literature, including the lack of a scalable, generative framework for autonomous policy creation in decentralized environments. By embedding policy evolution capabilities within edge AI, GPMs address the critical need for explainable, adaptive, and efficient governance mechanisms; this positions the current research at the frontier of integrating generative modeling with edge-centric autonomy.

3. Methodology

This study adopted a mixed-methods experimental design combining simulated edge AI environments and empirical evaluation of policy model performance. The central methodological framework uses a transformer-based policy generator to build on generative reinforcement learning.

- **Research Design:** Three edge scenarios were modeled: (i) a UAV swarm surveillance mission, (ii) real-time traffic signal control in a smart city setting, and (iii) resource allocation in a distributed IoT sensor network. Each scenario was simulated using the OpenAI Gym and customized edge RL environments with imposed constraints such as low compute availability, intermittent connectivity, and variable latency.
- **Data Collection:** Synthetic datasets were generated through simulations incorporating realistic operational metrics, such as drone battery levels, signal strengths, and bandwidth utilization. The initial training of generative policy agents used an offline dataset generated by expert heuristics. Subsequent learning phases relied on reinforcement signals (e.g., latency minimization, success rate maximization) without further supervision.
- **Model Implementation:** A lightweight transformer architecture was implemented to generate policy tokens conditioned on encoded environment states; the reward function was dynamically adjusted per task to reflect performance metrics and resource constraints. Furthermore, a multi-agent reinforcement learning approach

was used to simulate collaborative edge agents, while the comparative baselines included DQN, PPO, and rule-based systems.

- **Analytical Techniques:** The models were evaluated on convergence time, reward stability, policy novelty (measured by Levenshtein distance between generated policies), and resource efficiency. Statistical significance was verified using ANOVA and Wilcoxon tests, while ablation studies assessed the contribution of transformer layers, self-attention, and generative replay mechanisms.

4. Results

The generative policy models (GPMs) demonstrated superior adaptability and performance across all edge AI environments tested. In the UAV surveillance scenario, GPMs achieved a 22% increase in mission completion rates compared to PPO baselines while reducing average decision latency by 15 ms. In smart traffic control simulations, the GPMs reduced average vehicle wait time by 18% relative to static policies.

Resource utilization efficiency was markedly higher for GPMs due to their ability to anticipate and proactively reconfigure task strategies. Across all three simulations, the transformer-based policy generators maintained over 90% policy validity post-training and showed minimal degradation under constrained resource availability. Policy novelty scores indicated high generative diversity, enabling the agents to discover effective but non-obvious policy patterns.

The empirical evaluation confirmed hypotheses (1) and (2), as GPMs functioned autonomously without human supervision and evolved dynamically in response to environmental feedback. Hypothesis (3) was supported by statistically significant improvements in reward accumulation and policy robustness compared to traditional RL models.

5. Discussion

The findings align with prior assertions on the limitations of traditional RL in dynamic and resource-constrained edge settings (Yang et al., 2024; Khani et al., 2020). By embedding generative capabilities into the policy learning framework, GPMs introduce a novel paradigm for autonomous governance. The adaptability observed in GPM agents resonates with the theoretical expectations of generative policy spaces and supports broader application in real-time distributed systems.

This research contributes to the theoretical discourse on edge AI autonomy by illustrating how generative methods can self-produce diverse, context-sensitive policies. The use of transformers in generative RL proves particularly impactful due to their ability to retain long-range dependencies and synthesize complex, temporally coherent action sequences; this framework could underpin the next generation of autonomous edge systems, from disaster-response drones to intelligent building automation, by reducing reliance on human operators and central coordination. Moreover, it contributes to explainable AI by enabling retrospective inspection of policy generation pathways.

Research Limitations

Despite promising results, the study faced limitations that merit consideration. First, realistic simulations do not fully capture the unpredictability of real-world edge environments. Second, although optimized, the training and inference of transformer models still require a moderate computational budget, which may not be practical for ultra-low-power devices. Third, the generalizability of learned policies across drastically different task domains remains challenging despite observed adaptability within related contexts.

Future work should focus on testing GPMs in live edge deployments and further compressing model architectures for microcontroller-level feasibility. Moreover, there remains a need to explore ethical implications and fail-safe mechanisms in autonomous policy generation

6. Conclusion

This research establishes the feasibility and effectiveness of Generative Policy Models for autonomous governance in edge AI systems. However, GPMs enable self-evolving, adaptive, and efficient policy creation in decentralized environments through transformer-based generative reinforcement learning. The models outperform traditional baselines in decision-making speed, adaptability, and robustness; these contributions advance the field of edge AI by

offering a scalable framework for real-time, self-regulating systems operating under uncertainty and resource constraints.

6.1. Future Research

Building upon this foundation, future research could integrate federated generative policy learning to facilitate knowledge sharing across edge nodes while maintaining data privacy. Additionally, introducing continual learning mechanisms would allow edge agents to adapt policies over longer operational lifecycles; furthermore, hardware-specific optimizations, including neuromorphic implementations of GPMs, offer another promising direction. Finally, policy explainability and regulatory compliance frameworks must be developed to ensure trustworthy deployment in safety-critical applications such as healthcare robotics or autonomous transportation.

Compliance with ethical standards

Disclosure of conflict of interest

There is no conflict of interest to be disclosed.

References

- [1] Bengesi, S., El-Sayed, H., Sarker, M. K., Houkpati, Y., Irungu, J., and Oladunni, T. (2024). Advancements in Generative AI: A Comprehensive Review of GANs, GPT, Autoencoders, Diffusion Model, and Transformers. *IEEE Access*, 12, 69812–69837. <https://doi.org/10.1109/ACCESS.2024.3397775>
- [2] Bourechak, A., Zedadra, O., Kouahla, M. N., Guerrieri, A., Seridi, H., and Fortino, G. (2023). At the confluence of artificial intelligence and edge computing in IoT-based applications: A review and new perspectives. *Sensors*, 23(3), 1639–1688. <https://doi.org/10.3390/s23031639>
- [3] Cao, L. (2022). Decentralized AI: Edge intelligence and smart blockchain, metaverse, web3, and desci. *IEEE Intelligent Systems*, 37(3), 6–19. <https://doi.org/10.1109/MIS.2022.3181504>
- [4] Cao, Y., Sheng, Q. Z., McAuley, J., and Yao, L. (2023). Reinforcement learning for generative AI: A survey. *ArXiv*, 14(8), 1–30. <https://doi.org/10.48550/arXiv.2308.14328>
- [5] Chen, J., Ganguly, B., Xu, Y., Mei, Y., Lan, T., and Aggarwal, V. (2024). Deep generative models for offline policy learning: Tutorial, survey, and perspectives on future directions. *ArXiv*, 1–102. <https://doi.org/10.48550/arXiv.2402.13777>
- [6] Duan, S., Wang, D., Ren, J., Lyu, F., Zhang, Y., Wu, H., and Shen, X. (2022). Distributed artificial intelligence empowered by end-edge-cloud computing: A survey. *IEEE Communications Surveys and Tutorials*, 25(1), 591–624. <https://doi.org/10.1109/COMST.2022.3218527>
- [7] Golpayegani, F., Chen, N., Afraz, N., Gyamfi, E., Malekjafarian, A., Schäfer, D., and Krupitzer, C. (2024). Adaptation in edge computing: a review on design principles and research challenges. *ACM Transactions on Autonomous and Adaptive Systems*, 19(3), 1– 43. <https://doi.org/10.1145/3664200>
- [8] Hemmati, A., Raoufi, P., and Rahmani, A. M. (2024). Edge artificial intelligence for big data: a systematic review. *Neural Computing and Applications*, 36(19), 11461–11494. <https://doi.org/10.1007/s00521-024-09723-w>
- [9] Hu, X., Li, S., Huang, T., Tang, B., Huai, R., and Chen, L. (2023). How simulation helps autonomous driving: A survey of sim2real, digital twins, and parallel intelligence. *IEEE Transactions on Intelligent Vehicles*, 9(1), 593–612. <https://doi.org/10.1109/TIV.2023.3312777>
- [10] Hussain, F., Hussain, R., Hassan, S. A., and Hossain, E. (2020). Machine learning in IoT security: Current solutions and future challenges. *IEEE Communications Surveys and Tutorials*, 22(3), 1686–1721. <https://doi.org/10.1109/COMST.2020.2986444>
- [11] Khani, M., Sadr, M. M., and Jamali, S. (2024). Deep reinforcement learning-based resource allocation in multi-access edge computing. *Concurrency and Computation: Practice and Experience*, 36(15), e7995. <https://doi.org/10.1002/cpe.7995>
- [12] Kurt, G. K., Khoshkholgh, M. G., Alfattani, S., Ibrahim, A., Darwish, T. S., Alam, M. S., and Yongacoglu, A. (2021). A vision and framework for the high altitude platform station (HAPS) networks of the future. *IEEE Communications Surveys and Tutorials*, 23(2), 729–779. <https://doi.org/10.1109/COMST.2021.3066905>

- [13] Li, W., Hacid, H., Almazrouei, E., and Debbah, M. (2023). A comprehensive review and a taxonomy of edge machine learning: Requirements, paradigms, and techniques. *AI*, 4(3), 729–786. <https://doi.org/10.3390/ai4030039>
- [14] Rajapakse, V., Karunananayake, I., and Ahmed, N. (2023). Intelligence at the extreme edge: A survey on reformable TinyML. *ACM Computing Surveys*, 55(13s), 1–30. <https://doi.org/10.1145/3583683>
- [15] Shuvo, M. M. H., Islam, S. K., Cheng, J., and Morshed, B. I. (2022). Efficient acceleration of deep learning inference on resource-constrained edge devices: A review. *Proceedings of the IEEE*, 111(1), 42–91. <https://doi.org/10.1109/JPROC.2022.3226481>
- [16] Walia, G. K., Kumar, M., and Gill, S. S. (2023). AI-empowered fog/edge resource management for IoT applications: A comprehensive review, research challenges, and future perspectives. *IEEE Communications Surveys and Tutorials*, 26(1), 619–669. <https://doi.org/10.1109/COMST.2023.3338015>
- [17] Wu, J., Huang, C., Huang, H., Lv, C., Wang, Y., and Wang, F. Y. (2024). Recent advances in reinforcement learning-based autonomous driving behavior planning: A survey. *Transportation Research Part C: Emerging Technologies*, 164, 104654–104657. <https://doi.org/10.1016/j.trc.2024.104654>
- [18] Yang, N., Chen, S., Zhang, H., and Berry, R. (2024). Beyond the edge: An advanced exploration of reinforcement learning for mobile edge computing, its applications, and future research trajectories. *IEEE Communications Surveys and Tutorials*, 27(1), 546–594. <https://doi.org/10.1109/COMST.2024.3405075>
- [19] Zhu, J., Li, F., and Chen, J. (2024). A survey of blockchain, artificial intelligence, and edge computing for Web 3.0. *Computer Science Review*, 54, 100667–100670. <https://doi.org/10.1016/j.cosrev.2024.100667>