GAIA-DRL: A Geoenvironmental Agent for Energy Optimization in Batteryless IoT Networks with Ambient Backscatter

Edwardes A. Galhardo, Carlos B. Westphall, Wesley dos Reis Bezerra, and Antonio C. Oliveira Jr.

***Abstract*— This paper proposes GAIA-DRL, a framework that integrates ambient backscatter communication, geospatial intelligence, and deep reinforcement learning to optimize batteryless IoT networks. Real environmental data such as NDVI and pasture coverage are embedded into the agent's state vector, enabling adaptive, eco-aware communication control. Using the DDPG algorithm, GAIA-DRL jointly optimizes throughput, latency, energy efficiency, and interference under dynamic conditions. Simulations with real geospatial layers demonstrate improvements in performance and sustainability. The approach supports low-cost, low-carbon applications in 6G-ready IoT systems, including greenhouse gas monitoring, smart agriculture, and sustainable environmental management.**

***Keywords*— GAIA-DRL; Batteryless IoT; Ambient Backscatter; Geographic Information Systems (GIS); Remote Sensing; Deep Reinforcement Learning; Sustainable Networking.**

I. INTRODUCTION

The acceleration of climate change and its cascading effects on ecosystems and agriculture demand innovative solutions for real-time environmental monitoring. Advances in low-power communication and artificial intelligence have enabled the development of smart, energy-efficient systems capable of capturing ecological dynamics with minimal infrastructure.

Wireless Sensor Networks (WSNs) have long supported environmental monitoring [1], but their dependence on batteries introduces logistical and ecological limitations for large-scale or remote deployments. Ambient Backscatter Communication (AmBC) emerges as a promising alternative, enabling battery-free communication by harvesting and reflecting ambient radio frequency signals [2], [3].

To manage dense, dynamic, and energy-constrained networks, Deep Reinforcement Learning (DRL) provides a flexible approach. In particular, the Deep Deterministic Policy Gradient (DDPG) algorithm allows continuous control over system behavior, supporting fine-grained adjustments in transmission parameters based on real-time feedback [4], [5]. This is especially relevant in backscatter-based networks, where minor adjustments in reflectivity or timing can significantly affect performance under interference or mobility.

Edwardes A. Galhardo, Instituto de Informática, Universidade Federal de Goiás, Goiânia-GO, e-mail: edwardesamarogalhardo@inf.ufg.br; Carlos B. Westphall, Departamento de Informática e Estatística, Universidade Federal de Santa Catarina, Florianópolis-SC, e-mail: carlosbwestphall@gmail.com; Wesley dos Reis Bezerra, Universidade Federal de Santa Catarina, Florianópolis-SC, e-mail: wesley.bezerra@ifc.edu.br; Antonio C. Oliveira Jr, Instituto de Informática, Universidade Federal de Goiás, Goiânia-GO, e Fraunhofer Portugal AICOS, Porto, Portugal, e-mail: antoniojr@ufg.br.

Simultaneously, Geographic Information Systems (GIS) and remote sensing have become essential tools for tracking land-use changes, vegetation health, and environmental degradation. By combining in-situ IoT measurements with satellite-derived indices such as the Normalized Difference Vegetation Index (NDVI), researchers can build scalable and geospatially informed monitoring systems [6], [7].

In this work, we propose **GAIA-DRL**, a novel control algorithm for batteryless IoT networks. GAIA-DRL integrates spatial features such as NDVI and pasture coverage directly into the agent's state vector, enabling adaptive decision-making based on territorial and ecological priorities.

By combining AmBC, DRL, and geospatial intelligence, GAIA-DRL supports scalable, sustainable, and low-cost network architectures for applications in greenhouse gas tracking, precision agriculture, and smart environmental management.

II. RELATED WORK

The design of intelligent and sustainable IoT networks has attracted growing research interest, particularly in applications involving remote sensing, energy efficiency, and environmental monitoring.

Batteryless IoT and Ambient Backscatter: Ambient backscatter communication (AmBC) has emerged as a key enabler of batteryless IoT, allowing ultra-low-power devices to transmit by reflecting existing RF signals [2]. Early implementations demonstrated passive communication with minimal energy consumption [8], and recent advances extended AmBC to long-range protocols like LoRa [3]. The concept of Battery-Free IoT (BF-IoT) presents a promising alternative for sustainable deployments in remote areas [9].

Deep Reinforcement Learning in Network Optimization: DRL has been widely applied to optimize power control, scheduling, and data routing in wireless systems. Algorithms such as DDPG and PPO support continuous decision-making in dynamic environments [4]. In backscatter-based networks, DRL enables adaptive resource allocation under mobility and interference constraints [5].

GIS and Remote Sensing in Environmental Monitoring: Geospatial technologies are central to land-use mapping and ecosystem assessment. Platforms like Sentinel and CBERS provide vegetation indices (e.g., NDVI, EVI) that support pasture analysis, erosion tracking, and urban expansion monitoring [6], [7]. These variables enrich learning models with spatial context reflecting real environmental dynamics.

Hybrid IoT-GIS-DRL Approaches: Recent efforts have explored integrating IoT data, geospatial analysis, and AI. For instance, [10] used LULC projections for policy design in ecological corridors. However, few works incorporate satellite-derived variables directly into the control loop of IoT networks.

This paper introduces **GAIA-DRL**, a unified decision-making framework that merges AmBC, DRL, and geospatial data. Unlike prior studies that use spatial data mainly for monitoring, GAIA-DRL embeds environmental indicators such as NDVI and pasture coverage directly into the agent's state vector, enabling real-time, territory-adaptive communication strategies in dense, batteryless networks.

Inspired by recent approaches that emphasize adaptive architectures [11], GAIA-DRL advances this vision by incorporating geospatial intelligence into the decision logic of wireless agents, fostering intelligent behavior aligned with environmental conditions.

III. Proposed Method

This section describes the core components of the GAIA-DRL architecture, a novel approach for optimizing batteryless IoT networks by integrating ambient backscatter communication, deep reinforcement learning (DRL), and real geospatial data.

A. System Architecture

The proposed system consists of a dense IoT network composed of passive nodes spatially distributed across rural monitoring zones. These batteryless devices communicate by reflecting ambient RF signals—such as FM, AM, or TV—through ambient backscatter communication (AmBC). An intelligent controller based on DRL dynamically adjusts the reflective behavior of each node in real time to optimize communication while accounting for ecological constraints.

Figure 1 illustrates the GAIA-DRL architecture. It depicts an IoT-enabled pasture scenario where cows carry passive sensors that harvest and reflect ambient RF energy. The agent receives input from both communication metrics and environmental data—such as vegetation indices and pasture condition—obtained from remote sensing platforms like CBERS-4A and Sentinel satellites. These inputs form the agent's state vector, enabling environmentally adaptive control policies.

B. State Representation

At each time step t, the agent observes the environment via a multivariate state vector $\mathbf{s}_t$:

$$\mathbf{s}_t = [R_t, E_t, L_t, I_t, V_t]$$

Where:

- R_t = current throughput (kbps)
- E_t = estimated energy efficiency (a.u.)
- L_t = average latency (ms)
- I_t = interference level (a.u.)
- V_t = geospatial/environmental input (e.g., NDVI or pasture coverage)

The inclusion of V_t enables territory-aware and ecologically aligned decision-making.

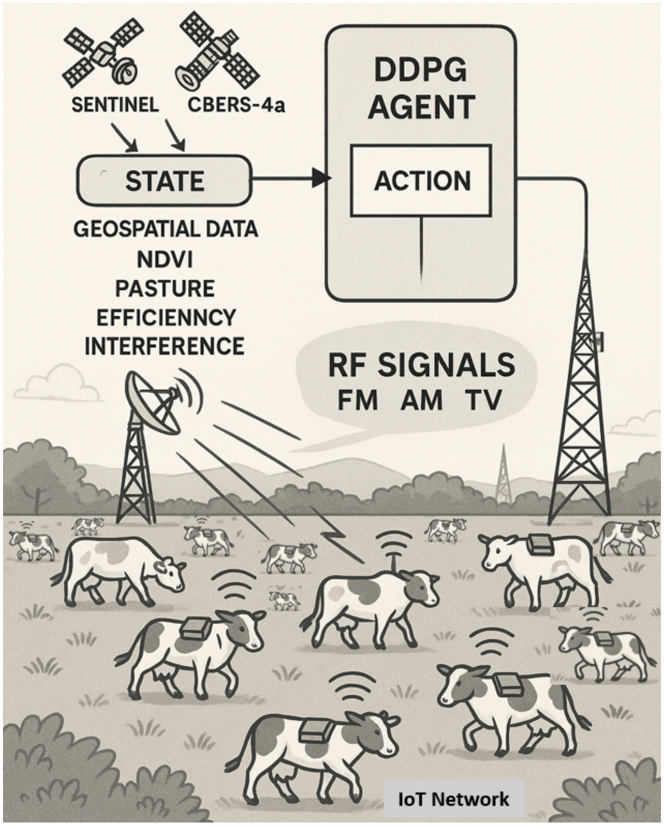

Fig. 1. GAIA-DRL architecture for adaptive control in batteryless IoT networks using ambient backscatter and geospatial intelligence.

C. Action and Control Mechanism

The agent outputs a continuous-valued action $a_t \in [0, 1]$ that modulates the logical reflection coefficient of each device. This controls the effective transmission strength, enabling dynamic trade-offs between communication performance and interference.

D. Reward Function

The agent is guided by a scalar reward r_t, computed using a weighted objective function:

$$r_t = \alpha_1 E_t + \alpha_2 R_t - \alpha_3 L_t - \alpha_4 I_t + \alpha_5 V_t$$

Weights $\alpha_1, ..., \alpha_5$ balance energy efficiency, throughput, latency, interference, and environmental context. In our setup:

$$[\alpha_1, \alpha_2, \alpha_3, \alpha_4, \alpha_5] = [0.25, 0.25, 0.2, 0.2, 0.1]$$

E. Learning Framework

We adopt the Deep Deterministic Policy Gradient (DDPG) algorithm [12], implemented in Python using `TensorFlow` and `Gym`. The agent is trained across multiple episodes in simulated environments with varying network densities and geospatial configurations. DDPG is well-suited for this scenario due to its support for continuous action spaces and ability to operate under non-stationary conditions with partial observability—characteristics inherent to backscatter-based IoT networks.

Key hyperparameters used in training include:

- Learning rate: 10^{-4} (actor), 10^{-3} (critic)
- Batch size: 64
- Replay buffer size: 10^5
- Discount factor γ: 0.99
- Soft update factor τ: 0.001

Training continues until the average reward stabilizes over a sliding window of 50 episodes.

F. GAIA-DRL Algorithm

Algorithm 1 outlines the full operation of the GAIA-DRL agent, including interaction with the environment, action selection, experience storage, and policy updates.

Algorithm 1 GAIA-DRL: DRL-Based Optimization Procedure

```
Initialize actor and critic networks with random weights
Initialize target networks and replay buffer
for each episode do
    Reset environment and get initial state s_0
    for each time step t do
        Select action a_t = μ(s_t) + N_t
        Execute a_t, observe r_t and next state s_{t+1}
        Store (s_t, a_t, r_t, s_{t+1}) in replay buffer
        Sample mini-batch and update networks via DDPG
    end for
end for
Return trained policy μ
```

IV. Experimental Setup

To evaluate the performance of the GAIA-DRL algorithm, we developed an experimental pipeline integrating network simulation, geospatial data processing, and deep reinforcement learning.

A. IoT Network Simulation

The wireless network was simulated using `OMNeT++`, modeling a dense mesh of passive sensor nodes that operate via ambient backscatter communication. The nodes were deployed with heterogeneous spacing and subject to terrain irregularities and obstacles to simulate realistic rural environments. Communication behavior was abstracted using logical reflection coefficients to emulate batteryless devices without active RF transmission.

We evaluated four different power levels (2, 5, 10, and 15 mW). For each level, two scenarios were compared:

- **Baseline:** static configuration with fixed reflection coefficients.
- **GAIA-DRL:** adaptive control using the DDPG agent.

Performance metrics collected include packet success rate, latency, interference, and energy efficiency. Each configuration was simulated over 100 runs to ensure statistical robustness. The mean and standard deviation of each metric were computed and reported in Section V.

B. Reinforcement Learning Environment

A custom OpenAI Gym environment was implemented in Python to simulate the interaction between the agent and the network. At each timestep, the agent receives a normalized state vector:

$$\mathbf{s}_t = [R_t, E_t, L_t, I_t, V_t]$$

The agent selects a continuous-valued action a_t that modulates the reflection behavior of each node. The environment returns updated metrics and a scalar reward based on the objective function defined in Section III.

Multiple network density scenarios and geospatial conditions were simulated to evaluate adaptability.

C. Integration with Geospatial Data

Geospatial layers were integrated using vegetation indices such as NDVI and pasture coverage maps, sourced from MapBiomas (Collection 8) and CBERS-4A satellite imagery. A preprocessing pipeline in Python (executed in Google Colab) extracted georeferenced control points and converted raster data into CSV format for use within the agent's training environment.

The resulting environmental variable V_t provided spatial awareness to the agent, enabling geospatially sensitive policy learning based on real territorial dynamics.

D. Pasture-Aware Monitoring Scenario

To simulate a relevant use case, we modeled a rural monitoring scenario in the municipality of Lages–SC, southern Brazil. NDVI data and pasture degradation maps identified zones with high ecological stress (e.g., overgrazing, low vegetation cover). The GAIA-DRL agent learned to prioritize these regions by increasing monitoring frequency and adjusting reflective behavior accordingly.

All results were visualized using `QGIS`, generating thematic maps that correlate network performance with environmental dynamics. All simulation scripts, geospatial layers, and training notebooks are available in our public repository: `https://github.com/edwardes-galhardo/GAIA-DRL1`, supporting full reproducibility.

V. Results and Discussion

This section presents the experimental results comparing the performance of the GAIA-DRL agent with a static baseline strategy. The evaluation focuses on throughput, latency, energy efficiency, interference, and spatial behavior of the network under varying transmission power levels. All metrics were averaged over 100 runs, and standard deviations are reported to assess variability.

A. Network Performance Optimization

Table I summarizes the success rate, energy efficiency, and latency across four power levels. The baseline refers to a fixed reflection strategy with no adaptivity. GAIA-DRL

TABLE I
SIMULATED NETWORK METRICS UNDER VARYING POWER LEVELS (MEAN ± STD)

Power (mW)	**Success Rate (%)** Baseline	GAIA-DRL	**Energy Efficiency** Baseline	GAIA-DRL
2	87.1 ± 1.9	94.6 ± 1.4	0.55 ± 0.03	0.68 ± 0.02
5	89.7 ± 1.6	95.7 ± 1.2	0.65 ± 0.02	0.77 ± 0.02
10	92.9 ± 1.2	96.8 ± 1.0	0.75 ± 0.01	0.82 ± 0.01
15	95.2 ± 1.0	97.1 ± 0.9	0.80 ± 0.01	0.85 ± 0.01

Power (mW)	**Latency (ms)** Baseline	GAIA-DRL
2	120 ± 5.4	105 ± 4.7
5	110 ± 4.9	95 ± 4.2
10	100 ± 4.3	85 ± 3.9
15	90 ± 3.8	75 ± 3.6

consistently outperformed the baseline, especially under low-power settings, improving communication reliability and responsiveness.

Figure 2 illustrates the gain in success rate across transmission levels. GAIA-DRL achieves up to 8% higher success in low-power regimes (2–5 mW) through adaptive tuning of node reflectivity.

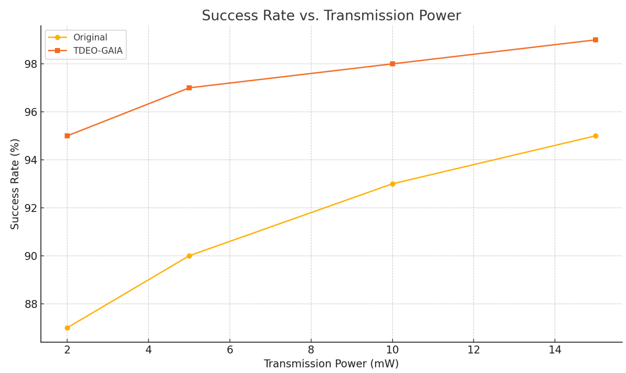

Fig. 2. Success rate vs. power level for baseline and GAIA-DRL.

Figure 3 shows the trade-off between energy efficiency and latency. GAIA-DRL reduced latency by up to 15% and increased efficiency by up to 25%, demonstrating its suitability for low-power, delay-sensitive IoT applications.

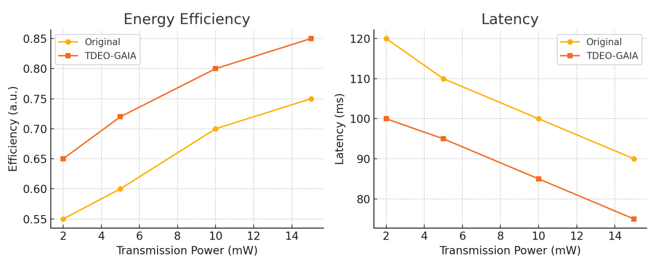

Fig. 3. Energy efficiency and latency across transmission levels.

B. Throughput and Interference Trade-Off

Table II presents throughput and interference metrics. GAIA-DRL consistently delivered higher throughput while reducing interference across all power levels.

TABLE II
THROUGHPUT AND INTERFERENCE COMPARISON (MEAN ± STD)

Power (mW)	**Throughput (kbps)** Baseline	GAIA-DRL	**Interference (a.u.)** Baseline	GAIA-DRL
2	12.3 ± 0.9	16.2 ± 1.0	0.25 ± 0.03	0.20 ± 0.02
5	18.1 ± 1.0	22.4 ± 1.1	0.35 ± 0.03	0.28 ± 0.02
10	24.2 ± 1.1	28.3 ± 1.0	0.50 ± 0.02	0.42 ± 0.02
15	26.6 ± 1.0	29.1 ± 0.9	0.65 ± 0.02	0.58 ± 0.01

C. Spatially Context-Aware Adaptation

D. Spatially Context-Aware Adaptation

By incorporating NDVI and pasture coverage data into the state vector, the GAIA-DRL agent was able to perceive spatial variations in vegetation health and land use intensity. This enabled the system to dynamically adapt its communication strategy based on ecological conditions observed across the monitored region.

Figure 4 presents a spatial transmission activity map for the municipality of Lages–SC. The underlying geospatial layer combines NDVI values and pasture degradation data from MapBiomas (Collection 8). In this representation, greener areas indicate healthy vegetation with minimal intervention needs, while lighter or yellowish regions correspond to ecologically stressed zones—such as overgrazed pastures or degraded land with low NDVI values.

The figure clearly demonstrates how the GAIA-DRL agent increases the frequency and intensity of transmissions in areas with poor vegetation conditions, focusing monitoring efforts where environmental risk is higher. Conversely, in stable areas with dense vegetation cover, the agent reduces transmission activity to conserve energy and avoid unnecessary interference.

This adaptive behavior exemplifies the system's ability to integrate geospatial intelligence into real-time network control, offering a sustainable approach to IoT monitoring that aligns communication decisions with territorial priorities and environmental vulnerability.

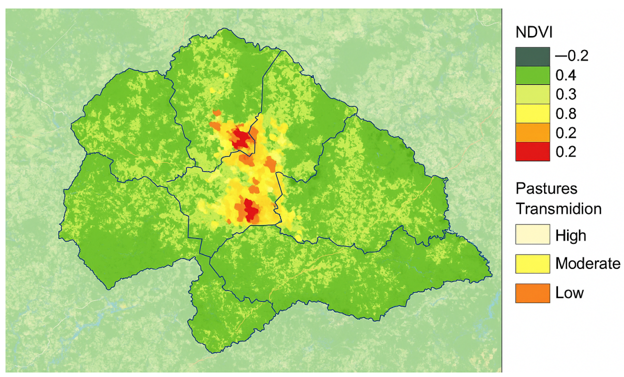

Fig. 4. Spatial distribution of transmission activity in Lages–SC, where zones with low NDVI and high pasture degradation (lighter areas) receive intensified monitoring by the GAIA-DRL agent, while stable vegetation zones (darker green) exhibit reduced transmission activity.

E. Critical Analysis and Limitations

Although results are promising, several limitations must be addressed. First, the current validation is based on simulations.

Field tests are required to assess real-world viability under interference, mobility, and irregular topography. Second, the current reward function uses fixed weights; future work should explore dynamic tuning strategies and multi-objective formulations.

Additional concerns include scalability to larger networks and robustness to outdated or inaccurate geospatial layers. Moreover, acquiring up-to-date NDVI or land-use data in real time may be constrained in regions with low satellite coverage or high cloud presence.

Despite these challenges, the results reinforce the feasibility and practical value of using DRL in batteryless, geospatially adaptive IoT networks for sustainable monitoring.

VI. Conclusion and Future Work

This paper presented **GAIA-DRL**, a reinforcement learning-based optimization framework for batteryless IoT networks operating with ambient backscatter communication. The proposed system integrates deep reinforcement learning (DDPG), real-time network metrics, and geospatial intelligence (NDVI and pasture coverage) to enable adaptive, context-aware communication control in dense and energy-constrained environments.

Simulation results demonstrated that GAIA-DRL improves success rate, energy efficiency, and latency, while reducing interference—even under low-power conditions. By embedding geospatial variables into the agent's state vector, the system adapts dynamically to ecological conditions, prioritizing monitoring in vulnerable or degraded areas. This capability is particularly relevant for sustainable agriculture, greenhouse gas tracking, and low-carbon 6G IoT infrastructures.

The integration of GIS and remote sensing expands the potential of reinforcement learning in real-world deployments, supporting scalable, cost-effective, and environmentally aware communication strategies.

Future work will focus on:

- Performing field validation with real AmBC devices and passive sensors, in collaboration with institutions such as EPAGRI/CIRAM;
- Developing a dynamic simulation platform with real-time geospatial data layers;
- Enhancing the reward function via adaptive weight tuning and multi-objective learning strategies sensitive to seasonal and spatial variations;
- Assessing scalability in larger deployments and improving robustness to outdated or incomplete environmental data;
- Exploring real-time NDVI and land-use data acquisition from alternative sources with higher temporal resolution.

We believe that combining batteryless IoT, artificial intelligence, and geospatial analysis offers a promising path toward resilient and ecologically responsible smart sensing infrastructures.

Acknowledgments

The authors would like to thank the Federal Institute of Tocantins (IFTO) for its institutional support and encouragement throughout the development of this research. They also extend their gratitude to the Agricultural Research and Rural Extension Company of Santa Catarina (EPAGRI) and the Foundation for the Support of Research and Innovation of the State of Santa Catarina (FAPESC) for their support, collaboration, and provision of geospatial and environmental datasets that significantly contributed to the development and validation of this work.

This study was partially funded by FAPESC through **Public Call FAPESC 57/2024**, in partnership with EPAGRI.

References

[1] X. Xu, Y.-J. Li, S. Zhou, and J. Wang, "A survey of energy harvesting wireless communications: Models and offline optimal policies," *IEEE Internet of Things Journal*, vol. 8, no. 8, pp. 6469–6489, 2021.

[2] J. Zhang, X. Lin, Y. Chen, and K. Wang, "Survey of ambient backscatter communication: Recent advances and future perspectives," *IEEE Internet of Things Journal*, vol. 9, no. 3, pp. 1869–1884, 2022.

[3] Z. Shen, Y. Xu, M. Li *et al.*, "Lora-based battery-free communication for sustainable iot," *IEEE Sensors Journal*, 2023.

[4] L. Zhao, X. Chen, and Y. Liu, "Deep reinforcement learning for resource optimization in wireless iot networks," *IEEE Transactions on Wireless Communications*, vol. 21, no. 5, pp. 3684–3699, 2022.

[5] L. Xie, J. Lin, Y. Zhang, and H. Wang, "Resource allocation in backscatter communication networks: A deep reinforcement learning approach," *IEEE Transactions on Cognitive Communications and Networking*, vol. 7, no. 3, pp. 935–949, 2021.

[6] A. Valjarevic *et al.*, "Gis and remote sensing techniques in sustainable land use and ecosystem management," *Land Use Policy*, vol. 113, p. 105899, 2022.

[7] M. Belgiu and O. Csillik, "Sentinel-2 cropland mapping using pixel-based and object-based time-weighted dynamic time warping analysis," *Remote Sensing of Environment*, vol. 204, pp. 509–523, 2018.

[8] V. Liu, A. Parks, V. Talla, S. Gollakota, D. Wetherall, and J. R. Smith, "Ambient backscatter: Wireless communication out of thin air," *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 4, pp. 39–50, 2013.

[9] W. Jiang, Y. Wang, L. Tang, and K. Wang, "Green iot: Battery-free technologies for sustainable sensing," *IEEE Internet of Things Journal*, 2023.

[10] J. Mio de Souza, P. Morgado, E. da Costa, and L. Vianna, "Predictive scenarios of lulc changes supporting public policies: The case of chapecó river ecological corridor, santa catarina/brazil," *Land*, vol. 12, no. 1, p. 181, 2023.

[11] A. L. de J. Gonçalves, L. A. Freitas, and A. Oliveira-Jr, "Arquitetura de dimensionamento adaptativo com suporte ao aprendizado," in *Anais do Simpósio Brasileiro de Redes de Computadores e Sistemas Distribuídos (SBRC)*, Brasil, 2025, disponível em: https://github.com/LABORA-INF-UFG/AdaptScale.

[12] J. Kim, G. Kim, S. Hong, and S. Cho, "Advancing multi-agent systems integrating federated learning with deep reinforcement learning: A survey," in *2024 Fifteenth International Conference on Ubiquitous and Future Networks (ICUFN)*, 2024, pp. 55–57.