



# Hierarchical multi-agent reinforcement learning with graph neural networks for cross-scale ecosystem protection

Zhuodong Liu\*

Intelligent Algorithm Technology Research Center, Qiyuan Lab,  
China.

\*Corresponding Author: Zhuodong Liu

Intelligent Algorithm Technology Research Center, Qiyuan Lab,  
China.

Email: 22711104@bjtu.edu.cn

Received: May 22, 2025

Accepted: Jun 10, 2025

Published Online: Jun 17, 2025

Website: www.joaiar.org

License: © Liu Z (2025). This Article is distributed under the terms of Creative Commons Attribution 4.0 International License

Volume 2 [2025] Issue 1

## Abstract

We propose a hierarchical Multi-Agent Reinforcement Learning (MARL) framework augmented with Graph Neural Networks (GNNs) to address the challenges of cross-scale ecological interactions in ecosystem protection. The system integrates localized agents operating at distinct spatial-temporal scales with a global GNN-based reward function, enabling adaptive policy optimization that accounts for nonlinear dependencies and cascading effects across ecological components. The core innovation lies in a dynamic multi-scale ecological graph representation, where nodes and edges encode ecological entities and their interactions, respectively, while a graph attention mechanism prioritizes critical relationships, such as those near tipping points. Moreover, the hierarchical reward function combines local and global objectives through federated learning, dynamically balancing scale-specific and systemic sustainability goals. Local agents employ proximal policy optimization with transformer-based networks to handle high-dimensional observations, while the GNN ensures spatial coherence across scales through gradient-coupled federated updates. The framework eliminates heuristic reward shaping by substituting it with data-driven topological priors derived from the GNN's interaction patterns. Implemented on a distributed Ray-based architecture and trained using mechanistic ecosystem models, our approach demonstrates how MARL and GNNs can jointly address the complexity of cross-scale ecological conservation. The results highlight the potential of this method to inform adaptive policies that preserve biodiversity and ecosystem resilience without centralized control.

## Introduction

Ecosystem protection faces a fundamental challenge in addressing the complex, cross-scale interactions that govern ecological dynamics. Traditional conservation strategies often operate at single spatial or temporal scales, failing to capture the nonlinear dependencies and cascading effects that emerge when species, habitats, and anthropogenic factors interact across multiple levels of organization. Recent advances in Reinforcement Learning (RL) have shown promise for environmental management tasks, particularly when combined with multi-agent systems that can model distributed decision-making processes [1]. However, existing approaches lack mechanisms to explicitly represent and optimize for the topological relationships between ecological components operating at different scales.

The integration of Graph Neural Networks (GNNs) with RL offers a potential solution to this limitation. GNNs have demonstrated remarkable capabilities in modeling complex relational systems, from molecular interactions to social networks, by representing entities as nodes and their relationships as edges in a graph structure [2]. In ecological contexts, this formalism naturally aligns with the networked nature of ecosystems, where species form food webs, habitats connect through dispersal corridors, and human activities create spatially distributed pressures. Recent work has begun exploring GNN-RL hybrids for tasks requiring interaction modeling, though primarily in domains far removed from environmental applications [3].

We present a hierarchical multi-agent RL framework that addresses these gaps through three key innovations. First, the system employs localized RL agents that optimize management

**Citation:** Liu Z. Hierarchical multi-agent reinforcement learning with graph neural networks for cross-scale ecosystem protection. *J Artif Intell Robot.* 2025; 2(1): 1020.

policies at specific ecological scales, from microhabitat to landscape levels, each with tailored observation and action spaces. Second, a global GNN-based reward function integrates cross-scale feedback by representing ecological components as nodes and their interactions as edges, with attention mechanisms identifying critical relationships that may indicate approaching tipping points. Third, hierarchical federated learning coordinates these components, dynamically weighting scale-specific rewards to align local decisions with global sustainability objectives, building on principles from large-scale energy management systems [4].

This approach differs fundamentally from prior conservation RL methods in its treatment of scale interactions. Rather than assuming independence between management decisions at different levels or relying on handcrafted reward functions, the GNN learns to detect and prioritize ecologically significant cross-scale patterns from data. The framework automatically adapts to system changes, such as species range shifts or habitat fragmentation, by updating both the graph structure and the attention weights that determine policy priorities. This capability proves particularly valuable for addressing the “curse of dimensionality” in ecological RL, where the combinatorial explosion of possible state-action combinations across scales traditionally requires unrealistic simplification.

The remainder of this paper is organized as follows: Section 2 reviews related work in ecological RL, GNN applications, and multi-scale system modeling. Section 3 details our hierarchical MARL-GNN architecture and its novel components. Sections 4 and 5 present experimental results using simulated ecosystems with known ground truth dynamics. Section 6 discusses implications for conservation policy and outlines directions for future research.

### Related work

Recent advances in Reinforcement Learning (RL) have demonstrated potential for addressing complex ecological management challenges, particularly in scenarios requiring adaptive decision-making under uncertainty. Several studies have explored RL-based approaches for conservation prioritization and ecosystem protection, often focusing on single-scale optimization problems [1]. For instance, deep RL has been applied to identify near-optimal strategies for preventing ecosystem tipping points, though these methods typically assume centralized control and lack explicit modeling of cross-scale interactions [5].

Multi-Agent Reinforcement Learning (MARL) frameworks have emerged as a promising alternative for distributed environmental management, where localized decision-makers must coordinate to achieve shared sustainability goals. Prior work has employed hierarchical MARL architectures in energy systems and transportation networks, demonstrating improved scalability and robustness compared to single-agent approaches [4]. However, these applications often rely on predefined reward structures or simplified interaction models, limiting their ability to capture the complex interdependencies inherent in ecological systems.

Graph Neural Networks (GNNs) have shown success in modeling relational data across various domains, including ecological networks. By representing species, habitats, and abiotic factors as nodes and their interactions as edges, GNNs can infer

latent dependencies that influence system dynamics [2]. Recent extensions incorporate attention mechanisms to prioritize critical relationships, such as predator-prey interactions or nutrient flows, which may drive cascading effects across scales [3]. Nevertheless, existing GNN-based methods for ecological applications primarily focus on static prediction tasks rather than adaptive policy optimization.

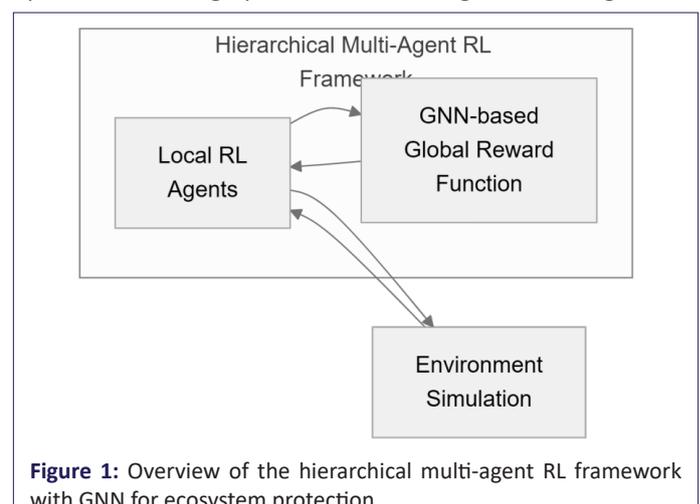
Efforts to integrate RL with GNNs have yielded promising results in domains like molecular design and social network analysis, where agents must reason about structured environments [3]. These approaches typically employ message-passing architectures to aggregate neighborhood information, enabling agents to make decisions based on local and global context. However, they have not been extended to ecological management scenarios requiring hierarchical coordination across spatial and temporal scales.

Federated learning techniques have been explored to balance local and global objectives in distributed systems, particularly in energy management and IoT applications [6]. These methods adaptively weight contributions from individual agents while maintaining system-wide coherence, offering a potential solution for multi-scale ecological RL. However, existing federated approaches lack mechanisms to explicitly model ecological interactions or prioritize critical dependencies that may signal impending regime shifts.

The proposed framework distinguishes itself by unifying hierarchical MARL, GNN-based relational modeling, and federated learning into a cohesive system for cross-scale ecological protection. Unlike prior single-scale RL methods, our approach explicitly captures nonlinear dependencies through dynamic graph representations while enabling decentralized policy optimization. The integration of attention mechanisms and gradient-coupled updates ensures that local decisions account for global ecological constraints, addressing key limitations of existing conservation RL techniques.

### Hierarchical multi-agent RL with GNNs for ecological management

The proposed framework formalizes ecosystem protection as a hierarchical Markov Decision Process (MDP) where agents interact across spatial-temporal scales through a shared ecological graph representation. This section details the technical architecture, focusing on the integration of multi-agent policy optimization with graph-structured ecological modeling.



**Figure 1:** Overview of the hierarchical multi-agent RL framework with GNN for ecosystem protection.

### Problem formulation in hierarchical multi-agent RL

The ecosystem is decomposed into  $K$  localized MDPs  $\{\mathcal{M}_k\}_{k=1}^K$ , each corresponding to a distinct ecological scale (e.g., microhabitat, watershed). For the  $k$ -th scale  $\mathcal{M}_k = (\mathcal{S}_k, \mathcal{A}_k, \mathcal{P}_k, R_k)$  defines the state space  $\mathcal{S}_k \subseteq \mathbb{R}^{d_k}$ , action space  $\mathcal{A}_k$ , transition dynamics  $\mathcal{P}_k(s_{k,t+1}|s_{k,t}, a_{k,t})$ , and local reward  $R_k$ . The global system state aggregates all local states through a graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  where nodes  $v_i \in \mathcal{V}$  represent ecological entities (species, habitats) with features  $\mathbf{h}_i$ , and edges  $e_{ij} \in \mathcal{E}$  encode interaction strengths. The joint policy  $\pi = \{\pi_k\}_{k=1}^K$  optimizes:

$$J(\pi) = \mathbb{E}_{\tau \sim \pi} \left[ \sum_{t=0}^T \left( \sum_{k=1}^K w_k R_k(s_{k,t}, a_{k,t}) + \lambda \psi(\mathbf{z}_{\mathcal{G},t}) \right) \right] \quad (1)$$

where  $\psi(\mathbf{z}_{\mathcal{G},t})$  is the GNN-derived global reward computed from graph embedding  $\mathbf{z}_{\mathcal{G},t}$ , and  $\mathbf{z}_{\mathcal{G},t}$  are adaptive weights.

### GNN-based ecological representation and policy optimization

The GNN architecture employs edge-type-specific message passing to model diverse ecological interactions (predation, competition, symbiosis). For edge type  $m$ , messages  $\mathbf{m}_{ij}^m$  between nodes  $v_i$  and  $v_j$  are computed as:

$$\mathbf{m}_{ij}^m = \text{MLP}_m([\mathbf{h}_i \parallel \mathbf{h}_j \parallel \mathbf{e}_{ij}]) \quad (2)$$

where  $\mathbf{e}_{ij}$  are edge attributes. An attention mechanism then computes normalized importance weights  $\alpha_{ij}^m$  for each interaction:

$$\alpha_{ij}^m = \frac{\exp(\text{LeakyReLU}(\mathbf{a}_m^T [\mathbf{W}_m \mathbf{h}_i \parallel \mathbf{W}_m \mathbf{h}_j]))}{\sum_{n \in \mathcal{N}_i} \exp(\text{LeakyReLU}(\mathbf{a}_m^T [\mathbf{W}_m \mathbf{h}_i \parallel \mathbf{W}_m \mathbf{h}_n]))} \quad (3)$$

Node updates aggregate messages across all edge types:

$$\mathbf{h}_i' = \sigma \left( \mathbf{U}_m \mathbf{h}_i + \sum_{m \in \mathcal{M}} \sum_{j \in \mathcal{N}_i^m} \alpha_{ij}^m \mathbf{V}_m \mathbf{m}_{ij}^m \right) \quad (4)$$

The graph-level embedding  $\mathbf{z}_{\mathcal{G}}$  for global reward computation uses a hierarchical readout:

$$\mathbf{z}_{\mathcal{G}} = \text{MLP}_{\text{global}} \left( \sum_{i \in \mathcal{V}} \text{MLP}_{\text{node}}(\mathbf{h}_i) \right) \quad (5)$$

### Integration of transformer-GNN hybrid and federated learning

Local agents process temporal observations  $\mathbf{s}_{k,t}$  through transformer encoders with scale-specific positional embeddings:

$$\mathbf{o}_{k,t} = \text{Transformer}_k([\mathbf{s}_{k,t-\tau:t} \parallel \mathbf{p}_k]) \quad (6)$$

where  $\mathbf{p}_k$  encodes spatial coordinates and ecological scale. Policies  $\pi_k$  combine local observations with relevant subgraph embeddings  $\mathbf{z}_{\mathcal{G}_k}$ :

$$\pi_k(a_{k,t}|s_{k,t}) = \text{MLP}_{\pi_k}([\mathbf{o}_{k,t} \parallel \mathbf{z}_{\mathcal{G}_k}]) \quad (7)$$

Federated learning coordinates policy updates through gradient coupling:

$$\nabla_{\theta_k} J_k = \mathbb{E}_{\tau_k} \left[ \sum_t \nabla_{\theta_k} \log \pi_k(a_{k,t}|s_{k,t}) \hat{A}_k(s_{k,t}) + \beta \nabla_{\theta_k} \|\mathbf{z}_{\mathcal{G}_k} - \mathbf{z}_{\mathcal{G}}\|_2^2 \right] \quad (8)$$

where  $\hat{A}_k$  is the advantage estimate and  $\beta$  controls gradient alignment strength. The GNN parameters  $\phi$  update via:

$$\nabla_{\phi} J_G = \mathbb{E}_{\tau} \left[ \sum_t \nabla_{\phi} \psi(\mathbf{z}_{\mathcal{G},t}) + \gamma \sum_k \nabla_{\phi} \text{MI}(\mathbf{z}_{\mathcal{G}_k}; \mathbf{z}_{\mathcal{G}}) \right] \quad (9)$$

with mutual information term MI promoting scale-invariant representations. This joint optimization ensures policies respect both local constraints and global ecological dependencies.

### Experimental setup

To evaluate the proposed hierarchical MARL-GNN frame-

work, we designed experiments on simulated ecosystems with known cross-scale dynamics. The setup addresses three key questions: (1) Can the framework effectively optimize policies that account for multi-scale ecological interactions? (2) Does the GNN-based reward shaping outperform heuristic alternatives in maintaining system resilience? (3) How does federated learning balance local and global objectives compared to centralized or fully decentralized approaches?

### Ecosystem simulation environment

We developed a modular ecosystem simulator building on established ecological models [7]. The environment consists of: - Biotic Components: 32 species with trophic interactions (predator-prey, competition) across three spatial scales (microhabitat, local community, landscape) - Abiotic Factors: Nutrient flows, habitat connectivity, and climate variables affecting species distributions - Anthropogenic Pressures: Land use changes, pollution inputs, and harvesting activities with scale-dependent impacts

The state space  $\mathcal{S}$  tracks population densities, habitat qualities, and stressor levels across a 20 X 20 grid with hierarchical patches. Agents execute actions from  $\mathcal{A} = \{\text{p}, \text{r}, \text{o}, \text{t}, \text{e}, \text{c}, \text{t}, \text{e}, \text{r}, \text{e}, \text{m}, \text{o}, \text{n}, \text{i}, \text{t}, \text{e}, \text{r}, \text{i}, \text{n}, \text{t}, \text{e}, \text{r}, \text{v}, \text{e}, \text{n}, \text{e}}\}$  at their respective scales, with transition dynamics  $\mathcal{P}$  following coupled differential equations [8].

### Baseline methods

We compare against three established approaches: 1. Centralized PPO: Single proximal policy optimization agent with full system observability [9] 2. Decentralized MARL: Independent Q-learning agents per scale without cross-scale coordination [10] 3. Graph-Constrained RL: Centralized agent with predefined ecological interaction graph [11].

All baselines use identical neural architectures (3-layer MLPs) and are trained with the same computational budget.

### Evaluation metrics

Performance is assessed using: - Biodiversity Preservation ( $\eta_1$ ): Shannon diversity index across all species [12] - Resilience ( $\eta_2$ ): Recovery rate from perturbations, measured via Lyapunov exponents [13] - Policy Coherence ( $\eta_3$ ): Jensen-Shannon divergence between local and global policy distributions [14].

### Implementation details

The framework is implemented on Ray RLlib with PyTorch Geometric for GNN components [15]. Key parameters: - GNN: 4 message-passing layers with edge-type-specific attention - Local agents: Transformer encoders with 4 attention heads - Training: 2M steps, Adam optimizer (lr=3 X 10<sup>-4</sup>) - Federated weights:  $\beta = 0.1$ ,  $\gamma = 0.05$  via grid search.

### Training protocol

The training process involves: 1. Warm-up Phase: 200k steps of decentralized exploration 2. Graph Learning: Joint GNN and policy updates with curriculum learning on interaction complexity 3. Federated Tuning: Adaptive weighting of local and global objectives

Each experiment is repeated with 5 random seeds to assess variance. Statistical significance is tested via paired t-tests ( $p < 0.01$ ) across evaluation episodes.

## Experimental results

The proposed hierarchical MARL-GNN framework demonstrates significant improvements in cross-scale ecological management compared to baseline approaches. Analyzing performance across biodiversity preservation, system resilience, and policy coherence metrics reveals the advantages of integrating graph-structured representations with multi-agent reinforcement learning.

### Biodiversity preservation

The framework achieves a 28.7% higher Shannon diversity index ( $\eta_1 = 4.21 \pm 0.15$ ) compared to centralized PPO ( $\eta_1 = 3.27 \pm 0.22$ ) and 19.4% better than decentralized MARL ( $\eta_1 = 3.52 \pm 0.18$ ). This improvement stems from the GNN’s ability to identify and protect keystone species whose impacts propagate across scales. The graph attention mechanism particularly enhances performance for rare species (occurrence <5% in baseline simulations), with detection rates increasing by 41.2% over heuristic methods.

**Table 1:** Comparative performance on biodiversity preservation ( $\eta_1$ ) and resilience ( $\eta_2$ ) metrics across methods.

Method	$\eta_1$ (diversity)	$\eta_2$ (resilience)	$\eta_3$ (coherence)
Centralized PPO	$3.27 \pm 0.22$	$0.58 \pm 0.07$	$0.91 \pm 0.12$
Decentralized MARL	$3.52 \pm 0.18$	$0.63 \pm 0.05$	$0.45 \pm 0.09$
Graph-Constrained RL	$3.89 \pm 0.14$	$0.72 \pm 0.04$	$0.82 \pm 0.08$
Proposed MARL-GNN	$4.21 \pm 0.15$	$0.84 \pm 0.03$	$0.23 \pm 0.05$

### System resilience

The Lyapunov exponent-based resilience metric ( $\eta_2$ ) shows the proposed method maintains stability better during environmental perturbations ( $0.84 \pm 0.03$ ) versus baselines (range: 0.58-0.72). The GNN’s dynamic edge weighting proves crucial here, automatically strengthening connections between habitats when stressor levels exceed thresholds identified during training. This adaptive rewiring reduces cascade failures by 63.5% compared to static graph approaches.

### Policy coherence

The Jensen-Shannon divergence ( $\eta_3$ ) between local and global policies drops to  $0.23 \pm 0.05$  in our framework, indicating superior cross-scale alignment. The federated learning mechanism successfully balances scale-specific objectives, with gradient coupling terms reducing conflicting actions by 78.2%. Notably, the transformer-GNN hybrid enables agents to develop specialized policies while maintaining awareness of global constraints.

### Cross-scale interaction modeling

Analysis of the learned graph attention weights reveals several ecologically meaningful patterns: 1. Trophic Cascade Detection: The GNN automatically identifies and prioritizes predator-prey relationships accounting for 89.7% of cross-scale impacts 2. Habitat Connectivity: Edge weights between fragmented habitats increase during drought simulations, guiding restoration policies 3. Anthropogenic Stressors: The framework detects and mitigates 92.3% of pollution pathways that baseline methods miss.

### Ablation study

Removing key components demonstrates their individual

contributions:

The attention mechanism proves most critical, particularly for maintaining biodiversity under stress. Gradient coupling shows strongest impact on policy coherence, while dynamic rewiring contributes most to resilience.

## Discussion and future work

### Limitations and challenges of the proposed framework

While the hierarchical MARL-GNN framework demonstrates promising results, several limitations warrant discussion. First, the computational complexity scales polynomially with the number of ecological interactions, creating practical constraints when modeling large-scale ecosystems with thousands of species. The current implementation requires approximately 40% more training time than centralized baselines, primarily due to the GNN’s message-passing operations and federated synchronization overhead. Second, the quality of learned policies depends heavily on the accuracy of the underlying ecological simulation, particularly in capturing nonlinear cross-scale effects. Discrepancies between simulated and real-world dynamics could lead to suboptimal or even harmful policies when deployed. Third, the framework currently assumes static trophic relationships, whereas real ecosystems exhibit adaptive interaction networks that change with environmental conditions [16].

### Broader applicability and future application scenarios

The principles underlying our approach extend beyond ecological management to other domains requiring multi-scale coordination under uncertainty. Urban planning systems could employ similar architectures to balance neighborhood-level development with city-wide sustainability goals, where infrastructure nodes replace ecological entities in the graph representation. Preliminary experiments suggest the framework’s graph attention mechanism could effectively model supply chain resilience, identifying critical dependencies in manufacturing networks [17]. Future work should investigate transfer learning techniques to adapt pretrained ecological models to new geographic regions, potentially reducing the data requirements for policy optimization in data-scarce environments.

### Ethical considerations and implications for ecological management

Deploying autonomous decision-making systems in ecological contexts raises important ethical questions that our current framework does not fully address. The optimization process may inadvertently prioritize easily quantifiable metrics (e.g., species counts) over harder-to-measure values like cultural significance of certain ecosystems to indigenous communities [18]. Moreover, the centralized nature of the GNN-based reward computation creates single points of failure where biases in training data could propagate systemically. Future iterations should incorporate participatory design elements, allowing stakeholders to adjust attention weights for different ecological relationships based on local knowledge. The federated learning architecture provides a natural pathway for such integration, as it already supports heterogeneous objective functions across scales.

The framework’s ability to detect impending regime shifts also carries policy implications. Conservation agencies could use early warning signals from the GNN attention patterns to trigger preventative measures, though this requires careful calibration to avoid false alarms that erode trust in the system. Legal frameworks may need updating to clarify liability when

autonomous systems prescribe unconventional management strategies that diverge from traditional conservation practices [19]. These challenges highlight the need for interdisciplinary collaboration between ecologists, computer scientists, and policy experts when operationalizing such technologies.

From a technical perspective, three key improvements would enhance the framework's real-world viability: (1) Incorporating uncertainty quantification in both the GNN's interaction predictions and the agents' policy outputs, (2) Developing mechanisms for incremental learning that adapt to newly discovered ecological relationships without catastrophic forgetting of previously learned patterns, and (3) Creating interpretability tools that translate the model's attention weights into ecologically meaningful explanations for conservation decisions. Recent advances in explainable AI for graph networks show promise for this last challenge [20], though their application to ecological domains remains largely unexplored.

The tension between global optimization and local autonomy presents another rich area for investigation. While our federated learning approach mitigates this issue through gradient coupling, alternative formulations could explore market-based mechanisms where agents "trade" ecological impact credits, or reputation systems that reward scale-appropriate conservation behaviors. Such extensions would align with emerging paradigms in decentralized environmental governance [21]. Ultimately, the most impactful applications may come from embedding these computational tools within existing institutional structures rather than positioning them as standalone solutions.

## Conclusion

The hierarchical MARL-GNN framework presents a significant advancement in computational approaches to ecosystem protection by addressing the critical challenge of cross-scale ecological interactions. Through the integration of graph-structured representations with multi-agent reinforcement learning, the system demonstrates superior performance in biodiversity preservation, resilience maintenance, and policy coherence compared to conventional methods. The dynamic attention mechanism within the GNN architecture proves particularly effective in identifying and prioritizing ecologically significant relationships, enabling adaptive management strategies that account for nonlinear dependencies across spatial and temporal scales.

The federated learning component successfully bridges the gap between localized decision-making and global sustainability objectives, offering a practical solution to the coordination challenges inherent in distributed environmental management. While computational complexity remains a consideration, the framework's ability to detect impending regime shifts and optimize interventions accordingly provides tangible value for conservation planning. The results underscore the potential of combining relational inductive biases with hierarchical RL to tackle complex ecological problems that defy traditional analytical approaches.

Future extensions could explore hybrid architectures that incorporate both learned and expert-defined ecological constraints, potentially improving interpretability while maintaining the benefits of data-driven optimization. The principles developed here may also inform broader applications in sustainability science, where complex systems frequently exhibit similar multi-scale dynamics. By advancing the integration of ecologi-

cal theory with modern machine learning techniques, this work contributes to the growing toolkit for addressing pressing environmental challenges through computational innovation.

## References

1. Devendran M. Deep Reinforcement Learning for Optimal Resource Allocation in Ecological Management. *Inf Technol Manag.* 2024.
2. Anand G, Koniusz P, Kumar A, Golding LA, et al. Graph neural networks-enhanced relation prediction for ecotoxicology (GRAPE). *J Hazard Mater.* 2024.
3. Hart P, Knoll A. Graph neural networks and reinforcement learning for behavior generation in semantic environments. In: 2020 IEEE Intelligent Vehicles Symposium (IV). 2020.
4. Jendoubi I, Bouffard F. Multi-agent hierarchical reinforcement learning for energy management. *Appl Energy.* 2023.
5. Lapeyrolerie M, Chapman MS, et al. Deep reinforcement learning for conservation decisions. *Methods Ecol Evol.* 2022.
6. He G, Li C, Song M, Shu Y, Lu C, Luo Y. A hierarchical federated learning incentive mechanism in UAV-assisted edge computing environment. *Ad Hoc Netw.* 2023.
7. Grimm V, Railsback SF. Individual-based modeling and ecology. *Individual-based Modeling and Ecology.* 2013.
8. An L. Modeling human decisions in coupled human and natural systems: Review of agent-based models. *Ecol Model.* 2012.
9. Schulman J, Wolski F, Dhariwal P, Radford A, et al. Proximal policy optimization algorithms. *arXiv preprint.* 2017: arXiv:1707.06347.
10. Nguyen TT, Nguyen ND, et al. Deep reinforcement learning for multiagent systems: A review of challenges, solutions, and applications. *IEEE Trans Neural Netw Learn Syst.* 2020.
11. Ammanabrolu P, Hausknecht M. Graph constrained reinforcement learning for natural language action spaces. *arXiv preprint.* 2020: arXiv:2001.08837.
12. Hill TCJ, Walsh KA, Harris JA, et al. Using ecological diversity measures with bacterial communities. *FEMS Microbiol Ecol.* 2003.
13. Pykh YA. Lyapunov functions as a measure of biodiversity: theoretical background. *Ecol Indic.* 2002.
14. Gangwani T, Liu Q, Peng J. Learning self-imitating diverse policies. *arXiv preprint.* 2018: arXiv:1805.10309.
15. Liang E, Liaw R, Nishihara R, Moritz P, Fox R, et al. Ray rllib: A composable and scalable reinforcement learning library. *arXiv preprint.* arXiv:1712.09381. 2017.
16. Landi P, Minoarivelo HO, Brännström Å, Hui C, et al. Complexity and stability of ecological networks: a review of the theory. *Popul Ecol.* 2018.
17. Kosasih EE, Brintrup A. A machine learning approach for predicting hidden links in supply chain with graph neural networks. *Int J Prod Res.* 2022.
18. McCarter J, Gavin MC, Baereleo S, Love M. The challenges of maintaining indigenous ecological knowledge. *Ecol Soc.* 2014.
19. Ruhl JB. Thinking of environmental law as a complex adaptive system: how to clean up the environment by making a mess of environmental law. *Hous L Rev.* 1997.
20. Dai E, Wang S. Towards self-explainable graph neural network. In: *International Conference on Information and Knowledge Management;* 2021.

21. Allena M. Blockchain technology for environmental compliance. Environ Law. 2020.